

IBM Maximo Asset Management
Version 7 Release 5

*Administering Maximo Asset
Management*



Note

Before using this information and the product it supports, read the information in “Notices” on page 445.

This edition applies to version 7, release 5, modification 0 of IBM Maximo Asset Management and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2008, 2012.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Chapter 1. Configuring the system . . . 1

Basic system configuration.	1
Configuring clustered systems	2
Performance improvements with clusters	2
Preparing to create clusters	5
Creating properties files for clusters	6
Configuring message-driven beans for clusters	7
Creating build files for clusters	10
Building Maximo EAR files for clusters	11
Building the RMI registry file	14
Creating and deploying clusters in WebSphere Application Server	14
Deploying the remote method invocation registry file in WebSphere Application Server	14
Creating clusters in WebSphere Application Server	16
Configuring Java Message Service for WebSphere Application Server	17
Deploying EAR files for clusters in WebSphere Application Server	25
Creating and deploying clusters in WebLogic Server	25
Deploying the remote method invocation registry file for WebLogic Server	26
Creating clusters in WebLogic Server	28
Configuring the Java Message Service for WebLogic Server.	29
Deploying EAR files for clusters in WebLogic Server	35
Building and deploying EAR files for basic configurations	36
Building EAR files for basic configurations	36
Deploying Maximo Enterprise Application Archive files in WebSphere Application Server.	37
Deploying Maximo Enterprise Application Archive files in WebLogic Server	38
Configuring general settings.	38
Online help configuration	38
Web application archive files	39
EAR files	39
Configuring application servers.	40
Memory settings for the application server process	40
Load balancing	41
Secure socket layer support	41
Creating Java virtual machines	41
Application server documentation	43
Configuring browser settings	43
Configuring Internet Explorer settings	43
Configuring session timeout periods	43
Migrating the administrative workstation	43

Chapter 2. Configuring databases . . . 45

Database design.	45
Relational database structure	45

Data dictionary tables	46
Integrity checker.	47
Storage partitions	47
Business objects	48
User-defined objects	49
Configuration levels for objects.	49
Database relationships.	52
Business object attributes	53
Attribute data types	53
Views	55
Indexes.	56
Defining objects for applications	56
Creating objects	56
Adding views to databases	57
Specifying attributes for objects.	59
Descriptions and long descriptions	59
Adding attributes to objects	59
Changing attributes.	60
Creating restrictions on attributes	61
Excluding user-defined attributes when duplicating objects	62
Enabling autonumbering for attributes	63
Adding tax types to database tables	63
Adding indexes	64
Creating relationships between parent and child tables	64
Account format configuration for General Ledger	65
General Ledger account configuration	65
General Ledger account components	66
General Ledger component sequence	66
General Ledger component lengths	67
Configuring the database.	67
Modes of configuring the database	67
Configuring the database in command-line mode	68
Configuring the database in administration mode	69
Restoring backup tables	70
Configuring the system for regulatory compliance	70
Electronic signatures and audit records	71
Login tracking	71
Electronic signature feature	72
Electronic audit records	72
Electronic signature authentication.	73
Creating a drop-down list for the Reason for Change field	74
Adding values to the Reason For Change domain	75
Database changes unrelated to eAudit	76
Database changes involving eAudit	77
Controlling changes to objects	78
Defining lookup maps.	78
Adding system messages.	79
Query definitions	80
Search option configuration for performance optimization	80
Text search function	80
Search type configuration.	80

Chapter 3. Configuring the system with multiple languages 85

Configuration of multiple languages overview.	85
Multiple language tables and associated columns	85
Multiple language utilities - translation data toolkit	86
Multiple languages and system table customizations	87
Multiple languages and translations	87
Enabling multiple language support	88
Enabling multiple languages on objects and attributes	88
Enabling attributes for multiple languages	88
Viewing characters from multiple languages	89
Localizing databases for unsupported base languages	89
Adding unsupported second languages to databases	90
Removing secondary languages from the database	91
Translating records through applications.	92
Setting languages for tooltips	92
Create a maxdemo database after installation	93

Chapter 4. Administering the database 95

Database administration overview.	95
Database backup and restoration	95
Types of backups	95
Offline and online backups	96
Database statistics updates	97
DBMS_STATS package.	97
SQL server update statistics	98
Database updates	98
Application patches	98
Database update for system options	98
UpdateDB and customer extensions	98
a_customer.xml file	99
product_description.xml file	99
Managing database administration	100
Updating the Maximo database	100
Running UpdateDB	100

Chapter 5. Optimizing system performance. 103

Database server performance	103
Optimization techniques for all databases	103
Database indexing.	104
Optimized access to data	105
Modifying sizes of sequence caches	105
Optimizing performance in DB2	107
Setting environment variables and registry variables for optimal performance	108
DB2 registry variables	109
Tuning database manager settings	110
DB2 database manager settings	110
Tuning database configuration settings	111
Enabling the REOPT(ONCE) bind option	112
DB2 database configuration settings	113
Reorganization of tables and indexes in DB2	116
Optimizing performance in Oracle Database	117

Oracle Database initialization parameters	118
IBM WebSphere Application Server performance tuning.	120
Thread pool sizes	121
Heap size values	122
Determining optimal heap sizes in WebSphere Application Server	123
JVM commands to optimize performance	124
HTTP server performance tuning.	125
IBM HTTP Server compression and load balancing.	126
Optimized settings for operating system configuration	127
Performance-related settings on AIX.	128
Performance-related network parameters for Windows and Red Hat Enterprise Linux	129
Developing performance tests	130
Determining test objectives.	130
Developing use cases.	131
Developing test strategies	131
Defining test environments.	132
Scenario: Developing performance tests to measure processor utilization	133

Chapter 6. Implementing security. . . . 139

Security Groups overview	139
Security groups and access to sites and applications	140
Types of security groups.	141
Security process	141
Authentication of users	141
Authorizations for security groups	143
Security profiles	146
Security profile of an organization with two security groups - example	147
Login tracking	149
Encryption and security	150
Hacking and denial-of-service attacks	151
Automatic creation of user records authenticated by LDAP	152
Combination of security groups	152
Combination of security groups - rules for data restrictions	152
Combination of security groups - rules for application authorization	153
Combination of security groups - rules for approval limits and tolerances.	153
Combination of security groups - rules for authorization of general ledger components	154
Combination of security groups - rules for labor authorization	154
Combination of security groups - rules for site authorization	155
Combination of security groups - rules for storeroom authorization	155
Application server security	155
Application server security - properties for user and group management	156
Security roles for the application server	156
Single sign-on environment for application server security	157

LDAP and application security servers	157
LDAP data synchronization	158
Synchronization of cron task parameters for application server security	159
Working with security groups	160
Adding security groups	160
Assigning start centers for security groups	161
Assigning sites to security groups	161
Adding users to security groups	161
Granting authorization privileges to security groups	163
Granting administrative login authorization for database configuration	163
Authorizing application privileges for security groups	164
Authorizing access to storerooms for security groups	165
Authorizing access to labor information for security groups	165
Authorizing security group access to general ledger components	166
Authorizing standard services for security groups	167
Overriding password duration for security groups	168
Specifying restrictions for security groups	168
Specifying data restrictions for security groups	168
Specifying collection restrictions for security groups	169
Specifying global data restrictions for security groups	170
Specifying purchasing limits and tolerances for security groups	171
Deleting users from security groups	171
Deleting security groups	172
Encrypting properties for security	173
Changing encrypted files for security	174
Configuring the system to use application server security	174
Configuring WebSphere Application Server for LDAP security	175
Configuring two directory servers	176
Configuring WebLogic Server for LDAP security	177
Changing cron task parameters for data synchronization	178
Activating cron tasks to synchronize data	179
Configuring WebSphere Application Server for incremental synchronization	179
Setting password requirements	180
Generating passwords	182
Enabling login tracking	183

Chapter 7. Registering users. 185

Users overview	185
Administrative users	185
Database users	186
System users	187
Configuration of self-registration for users	187
Self-registration for users	188

Security controls	189
Passwords for users	189
Password hints for users	190
Security authorizations for users	191
Security profiles for users	191
Database access for users	191
Default insert sites for users	192
User statuses	192
Working with users	193
Adding users	193
Assigning users to security groups	194
Authorizing users to assign other users to security groups	195
Changing persons associated with users	195
Changing the status of multiple users	196
Changing the status of users	196
Changing user settings	197
Changing user settings for inactive site access	197
Changing user settings for language, locale, and time zone	197
Changing user settings for screen readers	197
Changing user settings for storerooms and insert sites	197
Changing settings for storerooms and insert sites for multiple users	198
Changing general ledger accounts for users	198
Implementing security for users	199
Specifying passwords for new users	199
Changing system and database passwords for users	199
Specifying password hints for users	200
Specifying security groups for users	201
Specifying security profiles for users	201
Specifying security profiles for multiple users	202
Granting user access to Oracle and Structured Query Language server databases	203
Changing user access to Oracle and Structured Query Language server databases	203
Removing user access to Oracle and Structured Query Language server databases	204
Logging out and blocking users	204
Enabling login tracking	204
Setting user defaults	205
Copying users	206
Deleting users	207
Deleting security groups from user profiles	208

Chapter 8. Managing communication templates 209

Communications template overview	209
Communication templates and escalations	209
Communication templates and the service desk	209
Communication templates and workflow	209
Substitution variables for communication templates	210
Predefined communication templates	210
Recipients of communication templates	211
Attachments for communication templates	211
Communication logs	211
Working with communication templates	211

Creating communication templates	211
Adding email addresses as communication template recipients	213
Adding person groups as communication template recipients	213
Adding persons as communication template recipients.	214
Adding roles as communication template recipients.	215
Attaching documents to communication templates.	215
Attaching document folders to communication templates	215
Attaching files to communication templates	216
Attaching web pages to communication templates.	216
Linking records to communication templates	217
Copying communication templates	217
Changing communication templates.	218
Deleting communication templates	219
Changing the status of communication templates.	220

Chapter 9. Managing escalations . . . 223

Escalations overview	223
Escalation engine	223
Escalation logs	223
Structured Query Language Expression Builder	224
Escalation points	224
Predefined escalations	225
Escalations and service level agreements	227
Communication templates and notifications	227
Escalation record fields	228
Deletion rules for escalations	229
Working with escalations	229
Creating escalations	229
Defining escalation points	230
Validating escalations.	232
Activating escalations	233
Modifying escalations	234
Deactivating escalations	234

Chapter 10. Configuring e-mail listeners 237

Testing connectivity between the application server and mail server.	237
Syntax of the testemail command	238
E-mail Listeners overview	238
Email listeners components.	239
E-mail listeners process	239
Predefined workflow process for e-mail listeners	240
E-mail listeners definitions	240
Security settings for e-mail listeners	241
Communication templates for e-mail listeners	242
Preprocessors for e-mail listeners.	246
Object key delimiters.	247
Logging	248
Java Message Driven Bean	248

E-mail messages	248
Polling of mail servers for email messages	248
Status of e-mail records	250
E-mail attachments	251
Message thresholds	251
E-mail formats for e-mail listeners	253
Working with E-mail Listeners	259
Purging e-mail records from the staging table	259
Customizing the e-mail listener preprocessor	260
Changing the object key delimiter	260
Working with e-mail listeners definitions	261
Creating e-mail listener definitions	261
Deleting e-mail listener definitions	262
Configuring the queues for WebSphere Application Server	263
Adding servers to the Java Messaging Service bus for e-mail listeners	264
Creating the Java Messaging Service bus destination for the listener inbound queue.	265
Creating the Java Messaging Service connection factory.	266
Creating the listener inbound Java Messaging Service queue	266
Activating the listener inbound queue	267
Configuring the Message Driven Bean in WebSphere Application Server.	268
Configuring the Java Messaging Service queues for WebLogic Server	269
Adding file stores for e-mail listeners - WebLogic Server	270
Adding Java Messaging Service servers for e-mail listeners - WebLogic Server	271
Adding Java Messaging Service modules for e-mail listeners - WebLogic Server	272
Adding Java Messaging Service connection factories for e-mail listeners - WebLogic Server.	272
Adding Java Messaging Service queues for e-mail listeners - WebLogic Server	273
Activating Java Messaging Service connection factories for e-mail listeners - WebLogic Server.	273
Configuring the Message Driven Bean in WebLogic Server	274
Activating workflow processes for e-mail listeners	275
Configuring e-mail listeners to use Java Messaging Service queues	276
Creating communications for e-mail messages	277
Email listener does not process an email	278

Chapter 11. Managing cron tasks . . . 281

Cron task setup overview	281
Preexisting cron tasks	281
Access levels for cron tasks.	282
Cron task parameters.	283
Instances of cron tasks	283
Working with cron tasks.	284
Creating cron task definitions	284
Deleting cron task definitions	284
Working with instances of cron tasks	285

Creating cron task instances	285
Copying cron task instances	286
Changing cron task instances	286
Reloading cron task instances	286
Deleting cron task instances	287
Disabling cron tasks on an application server	287
Viewing hidden cron tasks	288
Chapter 12. Managing domains.	289
Domains overview	289
Applications associated with domains	289
Types of domains	290
ALN domains	290
Crossover domains	291
Numeric range domains	291
Synonym domains	292
TABLE domains	292
Foreign keys and TABLE domains	293
Domains and organizations or sites	293
Working with domains	294
Adding alphanumeric domains	294
Adding crossover domains	295
Adding numeric domains	297
Adding numeric range domains	298
Adding table domains	300
Associating domain values with conditions	302
Creating synonyms of internal values	303
Deleting synonyms of internal values	304
Deleting domains	305
Chapter 13. Configuring and administering attached documents	307
Attached Documents overview	307
Configuration of attached documents	307
Configuration of attached documents for a single computer	308
Configuration of attached documents for two computers and a local Hypertext Transfer Protocol server	308
Configuration of attached documents for two computers and a dedicated Hypertext Transfer Protocol server	309
Configuration of attached documents for multiple computers and multiple Hypertext Transfer Protocol servers	310
Alternative configurations for attached documents	312
Multi-purpose internet mail extension mappings for WebLogic Server	313
Working with attached documents	314
Managing document libraries	314
Adding file attachments to the library	314
Adding URLs to the library	314
Modifying existing documents	315
Attaching documents to records	315
Printing work packs in a UNIX environment	316
Maintaining document libraries	316
Adding document folders	316
Associating document folders with applications	317
Configuring attached documents in a single computer environment	318
Creating attached documents directories in a single-computer environment	318
Creating a Web application in a single-computer environment	319
Editing the httpd.conf file in a single-computer environment	320
Editing default file paths in System Properties in a single-computer environment .	321
Editing default file paths in System Properties for multiple computers and multiple Hypertext Transfer Protocol servers .	324
Changing paths for demo data library files in a single-computer environment	326
Configuring attached documents for two computers and a local Hypertext Transfer Protocol server	327
Creating attached documents directories for two computers and a local Hypertext Transfer Protocol server	327
Creating Web applications for two computers and a local Hypertext Transfer Protocol server	328
Editing default file paths in System Properties for two computers and a local Hypertext Transfer Protocol server	329
Editing default file paths in related applications for two computers and a local Hypertext Transfer Protocol server	331
Changing paths for demo data library files for two computers and a local Hypertext Transfer Protocol server	332
Configuring attached documents for two computers and a dedicated Hypertext Transfer Protocol server	333
Creating attached documents directories for two computers and a dedicated Hypertext Transfer Protocol server	333
Setting up the server for attached documents for two computers and a dedicated Hypertext Transfer Protocol server	334
Editing default file paths in System Properties for two computers and a dedicated HTTP server	335
Editing default file paths in related applications for two computers and a dedicated Hypertext Transfer Protocol server .	337
Changing paths for demo data library files for two computers and a dedicated Hypertext Transfer Protocol server	338
Configuring attached documents for multiple computers and multiple Hypertext Transfer Protocol servers	339
Creating attached documents directories for multiple computers and multiple Hypertext Transfer Protocol servers	339
Setting up the server for attached documents for multiple computers and multiple Hypertext Transfer Protocol servers	340

Editing default file paths in System Properties for multiple computers and multiple Hypertext Transfer Protocol servers	341
Editing default file paths in related applications for multiple computers and multiple Hypertext Transfer Protocol servers	343
Changing paths for demo data library files for multiple computers and multiple Hypertext Transfer Protocol servers	344

Chapter 14. Managing log files 347

Logging overview	347
Logging application components	347
Loggers	347
Appenders	347
Layouts	348
Loggers settings	348
Log file locations	348
Log file names	349
Loggers in multiple server environments	349
EventTracker filter	349
Working with logging	350
Creating logging.properties files	350
Specifying log file locations	351
Managing appenders	351
Automation scripts loggers	352
Cron task loggers	352
Escalation loggers	354
Integration framework loggers	354
Enabling the EventTracker filter	355
Logging events for specific applications or users	356
Stopping the logging of events	356
Log correlation	357

Chapter 15. Managing bulletin boards 359

Bulletin board overview	359
Communication logs for bulletin board messages	359
Working with bulletin boards	359
Viewing bulletin board messages	359
Creating bulletin board messages	360
Specifying audiences for bulletin board messages	360
Changing the status of bulletin board messages	360
Copying bulletin board messages	361
Viewing communication logs for bulletin board messages	361
Viewing the history of bulletin board messages	361
Deleting expired bulletin board messages	362

Chapter 16. Working with sets 363

Creating item sets or company sets	363
Changing item or company sets	364

Chapter 17. Managing organizations 365

Organizations overview	365
Application levels and data storage	365
Sites and organizations	366

Activation and deactivation of organizations and sites	366
Item sets	366
Autonumbering	367
ABC breakpoints and organizations	368
Enablement of repair facilities	369
Customization options for applications	369
Taxes for organizations	371
Drilldown options	372
Working with organizations	372
Creating organizations	372
Activating organizations	373
Deleting organizations	373
Clearing material reservations for work orders	373
Specifying options for work orders and ticket owners	374
Setting purchasing options	374
Associating properties with contracts for organizations	374
Associating terms and conditions with contracts for organizations	375
Specifying options for invoices	376
Specifying autonumbering for applications	377
Specifying autonumbering for special order items	377
Displaying user messages	378

Chapter 18. Managing calendars . . . 379

Calendars overview	379
Shift patterns for calendars	379
Exceptions to the standard calendar	379
Working with calendars	380
Creating calendars	380
Specifying shifts in calendars	380
Applying shifts to calendars	381
Specifying shift patterns in calendars	381
Copying calendars	382
Deleting calendars	382
Establishing work periods	383
Creating work periods	383
Changing work periods	383
Specifying non-working time for work periods	383

Chapter 19. Managing classifications 385

Classifications overview	385
Classification paths and hierarchies	385
Generate Description option	386
Associations of records with classifications	387
Classification searches	387
Actual and authorized configuration item classifications	388
Actual configuration items	388
Authorized configuration items	388
Actual configuration item classifications and authorized configuration item classifications	388
Attributes of classifications	389
Groupings of attributes	389
Apply Down Hierarchy option	389
Classifications planning	389

Types of records to be classified	389
Categories of items for reporting	390
Industry codes	390
Working with classifications	390
Creating classifications	390
Associating attributes with records	391
Modifying classifications	392
Modifying attributes	393
Adding attributes	393
Searching for classifications from application records	394
Defining associations between actual and authorized configuration item classifications	394
Adding units of measure	395
Configuring the signature option to enable classification images	396
Adding images to the asset topology view	396

Chapter 20. Managing charts of accounts 397

Chart of accounts overview.	397
General ledger account codes	397
Organizational default accounts for general ledgers	398
Merge of general ledger accounts.	398
Resource codes for general ledgers	398
Inactive component values	399
Working with chart of accounts	399
Working with general ledger accounts	399
Creating general ledger account codes	399
Creating general ledger component values	401
Changing component values in general ledger accounts.	401
Deleting general ledger component values	402
Changing general ledger account structures	403
Updating databases for general ledger accounts	404
Setting up accounts	404
Setting up organization default accounts	404
Specifying company-related accounts for general ledgers	405
Specifying external labor control accounts for general ledgers	405
Specifying financial periods for general ledgers	405
Closing financial periods	406
Specifying resource codes for general ledgers	406
Specifying validation options	407

Chapter 21. Working with cost management. 409

Creating cost management projects	409
---	-----

Chapter 22. Managing currency codes 411

Creating currency codes	411
Changing currency codes	411

Chapter 23. Setting system properties 413

Global properties	413
Instance properties	413
Options for system properties	413
System properties and encryption algorithms.	414
System properties that contain password information	415
Values of system properties in files and applications	415
Restoration of default values for system properties	416
Fetch stop limit memory errors	416

Chapter 24. System properties 419

Asset properties	419
Attached document properties.	419
Automation scripts properties	421
Bidirectional language properties.	422
Bulletin board property	422
Calendar property.	422
Classification item properties	422
Communication template property	423
Condition property	423
Cron task properties	423
Database properties	423
Deployed assets property	427
Email interaction system properties	428
Email listener properties.	429
Environment properties	429
E-signature properties	429
General ledger property	430
Inventory property	430
Issues and transfers property	430
Lightweight Directory Access Protocol integration properties	430
maximo.properties file	431
Migration Manager properties.	432
mxe.help properties	433
OSLC properties	434
Reorder property	434
Report integration properties	435
Security properties	437
Server properties	440
Utilities for logging and testing	442
Work order generation property	443
Workflow properties	444

Notices 445

Trademarks	446
----------------------	-----

Chapter 1. Configuring the system

Access to the business components and the web application are provided by an application server. A basic system configuration typically might support a user load of 50 users or less. A clustered configuration consists of clusters of Java virtual machines, can support a larger user load, and can scale up as the user load requirements increase.

Basic system configuration

A basic system configuration consists of a single instance of the system running on an application server. That server connects to a single instance of the database that is available on a database server.

If the integration framework is also configured for deployment, then you must set up additional messaging queues. The additional queues enable the system to send data to the external systems and receive data from the external systems by using queues.

The basic configuration is appropriate for the following situations:

- Development configuration
- Quality assurance configuration (to test the development work)
- Production system with a user load of 50 users or fewer users

A basic configuration might overload, depending on how much processing is performed within the application. If you need a configuration that handles more traffic than a basic configuration, then you can add Java virtual machines, or you can use the clustered configuration.

Even with fewer than 50 user loads, the basic system configuration can overload if there is significant processing. For example, scheduled jobs (such as cron tasks) and reports require significant memory and processing power. If the basic system configuration performs poorly, you can deploy the clustered configuration.

The default reporting engine is run from the application server that provides reporting capabilities.

The following diagram shows the main components in the basic configuration.




Figure 1. Basic system configuration

Configuring clustered systems

A typical deployment has four clusters: user interface, cron task, integration framework, and report. You must create copies of the properties files, message-driven bean files, build files, and EAR files, and then customize the files for each cluster. Then, you can create and deploy the clusters on your application server.

Related concepts:

Implementing high availability

Related information:

Access to services by inbound messages

Performance improvements with clusters

A cluster groups similar functions on two or more Java virtual machines (JVMs) to process a single function, such as scheduled cron tasks. Clusters connect to the same database but operate independently. For example, if the cron task cluster fails, users can still connect to the user interface cluster.

Users who access applications through a browser typically expect immediate responses from the server. A clustered configuration can be scaled to support more concurrent users with faster response times. For example, when the number of concurrent users increases, you can add more JVMs to the user interface cluster.

Resource-intensive operations, such as reports, cron tasks, and the integration framework can be configured to run in separate clusters. You can configure processes that do not require user interaction to run on virtual machines that are separate from the virtual machines that provide user interaction. For example, scheduled cron task jobs and inbound messages from external systems (integration framework) can each run on separate JVMs. If the system load requires more

resources, you can add JVMs to meet the increased need; hardware resource increase might be required as well. For example, if your integrated framework clustered environment routinely processes thousands of messages from external systems, you can add more JVMs to the cluster. The reliability of a system increases when the workload is distributed over multiple JVMs.

A typical deployment includes the following clusters:

User interface cluster

The user interface cluster is intended for users to access the system from a web browser.

Integration framework cluster

The integration framework cluster processes integration messages from message queues, and moves messages into the queues. This cluster uses Java Message Service (JMS), Hyper Text Transfer Protocol (HTTP) POST commands, web services, and Enterprise JavaBeans (EJB) technology.

Cron task cluster

The cron task cluster processes scheduled jobs. You can run scheduled jobs for integration tasks, escalations, Lightweight Directory Access Protocol (LDAP), or to run reports.

Report cluster

A dedicated reports cluster runs the Business Intelligence and Reporting Tools (BIRT) report engine. If you do not create a report cluster, then the BIRT report engine runs in each cluster, which can affect the performance of user interactive applications.

The process for creating a clustered environment involves the completion of tasks that are related to Maximo® Asset Management and followed by the completion of tasks that are related to the application server, which is either WebSphere® Application Server or WebLogic Server. The following table outlines the process of creating clusters:

Task	Purpose
Create a maximo.properties file for each cluster that you want to deploy.	You create separate properties files so that each cluster can have different settings. For example, you set properties to have all scheduled cron tasks run on the cron task cluster.
Create copies of the ejb-jar.xml file for each cluster that you want to deploy. If your deployment includes WebSphere Application Server, you also need to create and edit copies of the ibm-ejb-jar-bnd.xmi file. If your deployment includes WebLogic Server, you also need to create and edit copies of the weblogic-ejb-jar.xml file.	The ejb-jar.xml file and the ibm-ejb-jar-bnd.xmi file or the weblogic-ejb-jar.xml file are modified to configure message-driven beans for continuous queues.
Create copies of the buildmaximoear.cmd file for each cluster that you want to deploy.	The buildmaximoear.cmd files are used to create the individual EAR files for each cluster.
Build the EAR files.	The EAR files for each cluster are built based on the settings in the individual properties files, ejb-jar.xml files, and the ibm-ejb-jar-bnd.xmi or weblogic-ejb-jar.xml files.

Task	Purpose
Build the remote method invocation (RMI) registry file.	The rmireg.war file is used to create the RMI registry.
Deploy RMI.	RMI is deployed to the application server to create an independent RMI registry, which ensures that if a JVM fails, the RMI registry is still available.
Create the clusters.	The clusters are created by creating JVMs that are members of the cluster.
If you are setting up an environment with an integration framework cluster that is connected to an external system, configure the JMS.	JMS is used to communicate with external systems.
Deploy the EAR files for the clusters.	You deploy the EAR files on the application server so that each cluster supports its dedicated functions.

The following diagram shows an example of a clustered configuration that is integrated with an external system. The user interface cluster consists of an application server that has a BIRT report engine, a product instance, and online help. The user interface cluster is accessed by a web browser, which sends the requests through a web server load balancer. The integration cluster and cron task cluster each consist of a separate application server that has a BIRT report engine and a product instance. All three clusters connect to a single instance of the product database. The external integration consists of the external system, an inbound queue, and an outbound queue. Both queues receive messages from the integration cluster and the user interface cluster and send messages to the cron task cluster.




Figure 2. Example of a clustered configuration

Related concepts:

Implementing high availability

Related tasks:

“Creating clusters in WebLogic Server” on page 28

You can create as many clusters as your deployment requires.

“Creating clusters in WebSphere Application Server” on page 16

In the Integrated Solutions Console, you can create as many clusters as your deployment requires. Each cluster can consist of two or more Java virtual machine (JVM) cluster members.

Preparing to create clusters

Before you create clusters on the application server, you must first complete tasks on the Maximo Asset Management side. You create customized properties files for clusters, configure the message-driven bean files, create the build files, build the EAR files, and create the `rmireg.war` file.

Creating properties files for clusters

To separate tasks and functions between the clusters, you need to create, edit, and encrypt copies of the `maximo.properties` file. A typical clustered environment has separate clusters for processing user interface operations, cron jobs, reports, and integration transactions.

About this task

The `install_home` variable represents the installed location of the Maximo Asset Management folder, which by default is `ibm\SMP\maximo`.

The `maximo.properties` file has an encrypted password for the database user. An unencrypted version of the file, which is named `maximo.properties_orig`, is provided in the `install_home\etc\` directory. When you need to modify the `maximo.properties` file, you must use the unencrypted version.

To prepare to create clusters, you must first create a copy of the properties file for each cluster that you plan to deploy. In another step in the preparation, you create a build file for each cluster and you edit the build file to specify the name of the properties file for the cluster.

Procedure

1. Navigate to the `ibm\SMP\maximo` directory. Create a backup copy of the existing `maximo.properties` file, and then delete the existing `maximo.properties` file.
2. Create the copy of the properties file for the user interface cluster.
 - a. Copy the `ibm\SMP\etc\maximo.properties_orig` file to `install_home\applications\maximo\properties\maximo.properties`.
 - b. Open the `maximo.properties` file in a text editor, add the **donotrun** option, and list all of your cron tasks, except the `JMSQSEQCONSUMER` cron task, for example:

```
mxe.crontask.donotrun=JMSQSEQCONSUMER, BBCron,
ESCALATION, ESCESCBLTNEXP, REPORTLOCKRELEASE, REPORTLOCKRELEASE1,
REPORTUSAGECLEANUP, REPORTUSAGECLEANUP1
```
 - c. Add the line `mxe.report.birt.viewerurl=rpt_jvm_url` where `rpt_jvm_url` is the URL of the report cluster.
 - d. If reports are scheduled, ensure that the **mxe.report.birt.disablequeuemanager** option is set to 1. For example:

```
mxe.report.birt.disablequeuemanager = 1
```
 - e. Save and close the file.
 - f. At a command prompt, change to the `install_home\tools\maximo` directory and run `encryptproperties.bat`.
 - g. Rename the `maximo.properties` file so that it is identified with the user interface cluster, for example, `maximoui.properties`.
3. Create the copy of the properties file for the cron task cluster.
 - a. Copy the `ibm\SMP\etc\maximo.properties_orig` file to `install_home\applications\maximo\properties\maximo.properties`.
 - b. Open the `maximo.properties` file in a text editor and add the **donotrun** option for the `JMSQSEQCONSUMER` cron task, for example:

```
mxe.crontask.donotrun=JMSQSEQCONSUMER
```
 - c. Save and close the file.
 - d. At a command prompt, change to the `install_home\tools\maximo` directory and run `encryptproperties.bat`.

- e. Rename the `maximo.properties` file so that it is identified with the cron task cluster, for example, `maximocron.properties`.
4. Create the copy of the properties file for the integration framework cluster.
 - a. Copy the `ibm\SMP\etc\maximo.properties_orig` file to `install_home\applications\maximo\properties\maximo.properties`.
 - b. Open the `maximo.properties` file in a text editor, add the **donotrun** option and set the value to `all`, for example:


```
mxe.crontask.donotrun=ALL
```
 - c. Save and close the file.
 - d. At a command prompt, change to the `install_home\tools\maximo` directory and run `encryptproperties.bat`.
 - e. Rename the `maximo.properties` file so that it is identified with the integration framework cluster, for example, `maximomif.properties`.
5. Create the copy of the properties file for the report cluster.
 - a. Copy the `ibm\SMP\etc\maximo.properties_orig` file to `install_home\applications\maximo\properties\maximo.properties`.
 - b. Open the `maximo.properties` file in a text editor, add the **donotrun** option and set the value to `all`, for example:


```
mxe.crontask.donotrun=ALL
```
 - c. Add the **mxe.report.birt.disablequeuemanager** option, and set the value to `0`, for example:


```
mxe.report.birt.disablequeuemanager=0
```
 - d. Save and close the file.
 - e. At a command prompt, change to the `install_home\tools\maximo` directory and run `encryptproperties.bat`.
 - f. Rename the `maximo.properties` file so that it is identified with the report cluster, for example, `maximorpt.properties`.

Configuring message-driven beans for clusters

You need to create copies of the files that contain the code for the message-driven beans. You modify the files for the integration framework cluster to configure message-driven beans for continuous queues.

About this task

The `install_home` variable represents the installed location of the Maximo Asset Management folder, which by default is `ibm\SMP\maximo`.

Procedure

1. Create a copy of the `ejb-jar.xml` file for each cluster that you plan to deploy. For example, if you plan to deploy four clusters, create the following four copies:
 - `ejb-jarui.xml` for the user interface cluster
 - `ejb-jarcron.xml` for the cron task cluster
 - `ejb-jarmif.xml` for the integration cluster
 - `ejb-jarrpt.xml` for the reports cluster
2. Open the `ejb-jarmif.xml` file that you created for the integration framework cluster and uncomment the code for the following message-driven beans:
 - `MessageDriven_JMSContQueueProcessor_1`
 - `MessageDriven_JMSContQueueProcessor_2`

- JMSContQueueProcessor-1
- JMSContQueueProcessor-2

After you uncomment the section, the code in your file should match the following code:

```
<!-- MEA MDB -->
<message-driven id="MessageDriven_JMSContQueueProcessor_1">
  <ejb-name>JMSContQueueProcessor-1</ejb-name>
  <ejb-class>psdi.iface.jms.JMSContQueueProcessor</ejb-class>
  <transaction-type>Container</transaction-type>
  <message-destination-type>javax.jms.Queue</message-destination-type>
  <env-entry>
    <env-entry-name>MESSAGEPROCESSOR</env-entry-name>
    <env-entry-type>java.lang.String </env-entry-type>
    <env-entry-value>psdi.iface.jms.QueueToMaximoProcessor</env-entry-value>
  </env-entry>
</message-driven>

<!-- MEA MDB for error queue -->
<message-driven id="MessageDriven_JMSContQueueProcessor_2">
  <ejb-name>JMSContQueueProcessor-2</ejb-name>
  <ejb-class>psdi.iface.jms.JMSContQueueProcessor</ejb-class>
  <transaction-type>Container</transaction-type>
  <message-destination-type>javax.jms.Queue</message-destination-type>
  <env-entry>
    <env-entry-name>MESSAGEPROCESSOR</env-entry-name>
    <env-entry-type>java.lang.String </env-entry-type>
    <env-entry-value>psdi.iface.jms.QueueToMaximoProcessor</env-entry-value>
  </env-entry>
  <env-entry>
    <env-entry-name>MDBDELAY</env-entry-name>
    <env-entry-type>java.lang.Long </env-entry-type>
    <env-entry-value>30000</env-entry-value>
  </env-entry>
</message-driven>

<!-- MEA MDB -->
<container-transaction>
  <method>
    <ejb-name>JMSContQueueProcessor-1</ejb-name>
    <method-name>*</method-name>
  </method>
  <trans-attribute>Required</trans-attribute>
</container-transaction>

<!-- MEA MDB for error queue -->
<container-transaction>
  <method>
    <ejb-name>JMSContQueueProcessor-2</ejb-name>
    <method-name>*</method-name>
  </method>
  <trans-attribute>Required</trans-attribute>
</container-transaction>
```

3. If the application server for your deployment is WebSphere Application Server, create a copy of the *install_home*\applications\maximo\mboejb\ejbmodule\meta-inf\ibm-ejb-jar-bnd.xmi file for each cluster that you plan to deploy. For example, if you plan to deploy four clusters in WebSphere Application Server, create the following four copies:
 - ibm-ejb-jar-bndui.xmi for the user interface cluster
 - ibm-ejb-jar-bndcron.xmi for the cron task cluster
 - ibm-ejb-jar-bndmif.xmi for the integration framework cluster

- `ibm-ejb-jar-bndrpt.xml` for the reports cluster
4. If you are using WebSphere Application Server, open the `ibm-ejb-jar-bndmif.xml` file that you created for the integration framework cluster and uncomment the code for the following message-driven bean bindings:
- `ejbbnd:MessageDrivenBeanBinding`
 - `ejbbnd:MessageDrivenBeanBinding`

After you uncomment the bindings, the code in your file should match the following code:

```
<!-- MEA MDB -->
<ejbBindings xmi:type="ejbbnd:MessageDrivenBeanBinding"
xmi:id="MessageDrivenBeanBinding_1" activationSpecJndiName="intjmsact">
<enterpriseBean xmi:type="ejb:MessageDriven"
href="META-INF/ejb-jar.xml#MessageDriven_JMSContQueueProcessor_1"/>
</ejbBindings>

<!-- MEA MDB for error queue -->
<ejbBindings xmi:type="ejbbnd:MessageDrivenBeanBinding"
xmi:id="MessageDrivenBeanBinding_1" activationSpecJndiName="intjmsacterr">
<enterpriseBean xmi:type="ejb:MessageDriven"
href="META-INF/ejb-jar.xml#MessageDriven_JMSContQueueProcessor_2"/>
</ejbBindings>
```

5. If the application server for your deployment is WebLogic Server, create four copies of `install_home\applications\maximo\mboejb\ejbmodule\meta-inf\weblogic-ejb-jar.xml`. For example, if you plan to deploy four clusters on WebLogic Server, create the following four copies:
- `weblogic-ejb-jarui.xml` for the user interface cluster
 - `weblogic-ejb-jarcron.xml` for the cron task cluster
 - `weblogic-ejb-jarmif.xml` for the integration cluster
 - `weblogic-ejb-jarrpt.xml` for the reports cluster
6. If you are using WebLogic Server, open the `weblogic-ejb-jarmif.xml` file that you created for the integration framework cluster and uncomment the code for the `JMSContQueueProcessor` sections.

After you uncomment the section, the code in your file should match the following code:

```
<!-- MEA MDB-->
<weblogic-enterprise-bean>
<ejb-name>JMSContQueueProcessor-1</ejb-name>
<message-driven-descriptor>
<pool>
<max-beans-in-free-pool>3</max-beans-in-free-pool>
</pool>
<destination-jndi-name>jms/maximo/int/queues/cqin</destination-jndi-name>
<connection-factory-jndi-name>jms/maximo/int/cf/intcf
</connection-factory-jndi-name>
</message-driven-descriptor>
<transaction-descriptor>
<trans-timeout-seconds>600</trans-timeout-seconds>
</transaction-descriptor>
<jndi-name>JMSContQueueProcessor-1</jndi-name>
</weblogic-enterprise-bean>

<weblogic-enterprise-bean>
<ejb-name>JMSContQueueProcessor-2</ejb-name>
<message-driven-descriptor>
<pool>
<max-beans-in-free-pool>3</max-beans-in-free-pool>
</pool>
<destination-jndi-name>jms/maximo/int/queues/cqinerr</destination-jndi-name>
```

```

<connection-factory-jndi-name>jms/maximo/int/cf/intcf
</ connection-factory-jndi-name>
</message-driven-descriptor>
<transaction-descriptor>
  <trans-timeout-seconds>600</trans-timeout-seconds>
</transaction-descriptor>
<jndi-name>JMSContQueueProcessor-2</jndi-name>
</weblogic-enterprise-bean>

```

Related concepts:

Integration framework overview

Creating build files for clusters

You must create a separate `buildmaximoear.cmd` file for each cluster. When you run the separate `buildmaximoear.cmd` files, you create a separate EAR file for each cluster.

About this task

The `install_home` variable represents the installed location of the Maximo Asset Management folder, which by default is `ibm\SMP\maximo`.

Procedure

1. Create a copy of the `install_home\deployment\buildmaximoear.cmd` file for each cluster that you plan to deploy. For example, if you plan to deploy four clusters, create the following four copies:
 - `buildmaximoearui.cmd`
 - `buildmaximoearcron.cmd`
 - `buildmaximoearmiif.cmd`
 - `buildmaximoearrpt.cmd`
2. Open the `buildmaximoear.cmd` file for each cluster, and add the following shell command at the beginning of the file where `cluster` is the cluster whose file you are editing:

```

copy /Y install_home\applications\maximo\properties\
maximocluster.properties
install_home\applications\maximo\properties\maximo.properties

```

For example, for the user interface cluster, you add the following shell commands:

```

copy /Y install_home\applications\maximo\properties\maximoui.properties
install_home\applications\maximo\properties\maximo.properties

```

3. In the `buildmaximoear.cmd` file for each cluster, add the following shell command at the beginning of the file where `cluster` is the cluster whose file you are editing:

```

copy /Y install_home\applications\maximo\mboejb\ejbmodule\meta-inf\ejb-
jarcluster.xml
install_home\applications\maximo\mboejb\ejbmodule\meta-inf\ejb-jar.xml

```

For example, for the user interface cluster, you add the following shell commands:

```

copy /Y install_home\applications\maximo\mboejb\ejbmodule\meta-inf\ejb-
jarui.xml
install_home\applications\maximo\mboejb\ejbmodule\meta-inf\ejb-jar.xml

```

4. If the application server for your deployment is WebSphere Application Server, in the `buildmaximoear.cmd` file for each cluster, add the following shell command at the beginning of the file where *cluster* is the cluster whose file you are editing:

```
copy /Y install_home\applications\maximo\mboejb\ejbmodule\meta-inf\ibm-  
ejb-jar-bndcluster.xml  
install_home\applications\maximo\mboejb\ejbmodule\meta-inf\ibm-ejb-jar-  
bnd.xml
```

For example, for the user interface cluster, you add the following shell commands:

```
copy /Y install_home\applications\maximo\mboejb\ejbmodule\meta-inf\ibm-  
ejb-jar-bndui.xml  
install_home\applications\maximo\mboejb\ejbmodule\meta-inf\ibm-ejb-jar-  
bnd.xml
```

5. If the application server for your deployment is WebLogic Server, in the `buildmaximoear.cmd` file for each cluster, add the following shell command at the beginning of the file where *cluster* is the cluster whose file you are editing:

```
copy /Y install_home\applications\maximo\mboejb\ejbmodule\meta-inf\  
weblogic-ejb-jarcluster.xml  
install_home\applications\maximo\mboejb\ejbmodule\meta-inf\weblogic-ejb-  
jar.xml
```

For example, for the user interface cluster, you add the following shell commands:

```
copy /Y install_home\applications\maximo\mboejb\ejbmodule\meta-inf\  
weblogic-ejb-jarui.xml  
install_home\applications\maximo\mboejb\ejbmodule\meta-inf\weblogic-ejb-  
jar.xml
```

6. For each cluster that you plan to deploy, in the `buildmaximoear.cmd` file, set `EAR_FILENAME` to the unique file name for the cluster. For example, if you are editing the `buildmaximoearui.cmd` file for the user interface cluster, set the file name to `maximoui.ear`.

```
set EAR_FILENAME=maximoui.ear
```

When you run the build script to build the EAR file, the resulting EAR file is named `maximoui.ear`.

Related concepts:

“EAR files” on page 12

EAR files are archives that contain all the required files to run an application.

“Web application archive files” on page 13

Web application archive (WAR) files are part of EAR files. They contain, for example, JSP or HTML pages.

Building Maximo EAR files for clusters

After you create a build file for each cluster, you must build a Maximo EAR file for the cluster. The name of the EAR file is based on the `set EAR_FILENAME` statement in the build file.

About this task

The `install_home` variable represents the installed location of the Maximo Asset Management folder, which by default is `ibm\SMP\maximo`.

Procedure

1. From the command prompt, navigate to *install_home*\maximo\deployment\ and run each of the four build files that you created for the clusters, for example:

- buildmaximoealui.cmd
- buildmaximoearcron.cmd
- buildmaximoearmif.cmd
- buildmaximoearrpt.cmd

Each build file creates a separate Maximo EAR file for the cluster.

2. Navigate to *install_home*\maximo\deployment\ and run **buildmxiehsear.cmd**.

Related concepts:

“EAR files”

EAR files are archives that contain all the required files to run an application.

“EAR files”

EAR files are archives that contain all the required files to run an application.

“Web application archive files” on page 13

Web application archive (WAR) files are part of EAR files. They contain, for example, JSP or HTML pages.

Related tasks:

“Building and deploying EAR files for basic configurations” on page 36

You can build and deploy EAR files for a basic configuration. In a clustered configuration, each cluster has its own EAR file to build and deploy.

“Building EAR files for basic configurations” on page 36

The EAR file contains all the fields required to run an application. There are two EAR files: maximo.ear and maximoiehs.ear. Both the EAR files contain one or more web application modules. In a clustered configuration, each cluster requires a separate build file to create an EAR file for the cluster.

EAR files:

EAR files are archives that contain all the required files to run an application.

The following two EAR files are used. Each EAR file contains one or more web application modules (.war extension):

- maximo.ear
 - maximouiweb.war
 - mboweb.war
 - meaweb.war
- maximoiehs.ear
 - iehs.war

You rebuild and redeploy EAR files whenever you:

- Modify .xml files or custom class files (maximo.ear).
- Modify HTML help topics (online help) (maximoiehs.ear).
- Modify settings in the maximo.properties file (Maximo.ear).

Related concepts:

“Web application archive files”

Web application archive (WAR) files are part of EAR files. They contain, for example, JSP or HTML pages.

Related tasks:

“Building EAR files for basic configurations” on page 36

The EAR files contains all the fields required to run an application. There are two EAR files: maximo.ear and maximoiehs.ear. Both the EAR files contain one or more web application modules. In a clustered configuration, each cluster requires a separate build file to create an EAR file for the cluster.

“Creating build files for clusters” on page 10

You must create a separate buildmaximoeear.cmd file for each cluster. When you run the separate buildmaximoeear.cmd files, you create a separate EAR file for each cluster.

“Building Maximo EAR files for clusters” on page 11

After you create a build file for each cluster, you must build a Maximo EAR file for the cluster. The name of the EAR file is based on the set `EAR_FILENAME` statement in the build file.

Web application archive files:

Web application archive (WAR) files are part of EAR files. They contain, for example, JSP or HTML pages.

WAR file	Description
maximouiweb.war	Contains the user interface-related JavaServer Pages (.jsp files), Java classes, static HTML files, and static image files. The buildmaximoeear.xml file has information about the files in this module. This web application uses the configuration details in the web.xml file, located in the <maximo root>\applications\Maximo\Maximouiweb\webmodule\WEB-INF folder. This file also specifies the URL to access online help.
mboweb.war	Contains the business objects, Java classes, and dependent third-party Java classes.
meaweb.war	The integration framework enables the exchange of application data with another application or with an external system. Users can create and maintain data in one system and use the integration framework to transfer data to an external system, which eliminates duplicate processing.
iehs.war	Provides the online help. The buildmxiehsear.xml file has information about all the files in this module.
rmireg.war	Creates the remote method invocation (RMI) registry file.

Related concepts:

“EAR files” on page 12

EAR files are archives that contain all the required files to run an application.

Related tasks:

“Creating build files for clusters” on page 10

You must create a separate `buildmaximoear.cmd` file for each cluster. When you run the separate `buildmaximoear.cmd` files, you create a separate EAR file for each cluster.

“Building Maximo EAR files for clusters” on page 11

After you create a build file for each cluster, you must build a Maximo EAR file for the cluster. The name of the EAR file is based on the set `EAR_FILENAME` statement in the build file.

Building the RMI registry file

The `rmi reg.war` file is used to create the remote method invocation (RMI) registry. After you create the `rmi reg.war` file, you can deploy the file on the application server.

About this task

The `install_home` variable represents the installed location of the Maximo Asset Management folder, which by default is `ibm\SMP\maximo`.

Procedure

From a command prompt, navigate to the `install_home/deployment` directory, and then run the following command: `buildrmi regwar.cmd`

Related concepts:

“Java remote method invocation” on page 15

Remote method invocation (RMI) is an application programming interface that provides a way for objects in separate memory areas to interact. Separate memory areas can be part of the same physical system or can be on different systems connected by a network.

Creating and deploying clusters in WebSphere Application Server

In WebSphere Application Server, you deploy the remote method invocation (RMI), create the clusters, configure JMS, and deploy the EAR files.

Related concepts:

“Application server documentation” on page 43

For more information about your application server, see the following web sites.

Deploying the remote method invocation registry file in WebSphere Application Server

A server that has a remote method invocation (RMI) registry continues to run even if another server in the cluster fails. When you build the RMI registry file, the file must be deployed in the application server.

About this task

For each physical server or virtual machine that acts as a server in your environment, you must create a JVM to host the RMI registry. The RMI JVM must be created outside of the clusters and must be started before any other JVMs.

The *install_home* variable represents the installed location of the product folder.

Procedure

1. Log in as the administrative user to the Integrated Solutions Console by using the address `http://servername:9060/admin`
2. In the navigation pane of the Integrated Solutions Console, select **Servers > Server Types > WebSphere Application Servers** and then click **New** to create a JVM for the RMI registry.
3. Specify **RMIRegistry** for the server name and use the default settings for the server template and server properties.
4. Click **Finish**, and then save the server information.
5. In the navigation pane of the Integrated Solutions Console, click **Applications > Application Types > WebSphere Enterprise Applications**.
6. Click **WebSphere enterprise applications**, and then click **Install**.
7. Click **Browse** and specify the *install_home*/deployment/default/rmireg.war file.
8. In the **Context Root** field, specify **RMI**.
9. Select the **Generate Default Bindings** check box.
10. Continue to the next section without changing any selections. Do not change the default application name `rmireg_war`.
11. In the Mapping Modules to Servers section, select the **RMIRegistry** server, select the **MBO Web Application** check box, and click **Apply**.
12. Confirm that the MBO Web Application is mapped to the **RMIRegistry**.
13. Click **Finish** to complete the installation of the file.
14. Save the `rmireg_war` file to the master configuration.
15. In the navigation pane of the Integrated Solutions Console, click **Applications > Application Types > WebSphere application servers**, and click the link for the `rmireg_war` application.
16. Repeat steps 1-14 for every physical server or virtual machine that acts as a server.

Related concepts:

“Java remote method invocation”

Remote method invocation (RMI) is an application programming interface that provides a way for objects in separate memory areas to interact. Separate memory areas can be part of the same physical system or can be on different systems connected by a network.

Java remote method invocation:

Remote method invocation (RMI) is an application programming interface that provides a way for objects in separate memory areas to interact. Separate memory areas can be part of the same physical system or can be on different systems connected by a network.

An RMI Registry is an area in memory that maintains the RMI address information of a Java object server. By default, the RMI registry is created on port 1099. More than one RMI registry can exist in memory. Each registry has a designated TCP/IP port for access.

If the `mxe.allowLocalObjects` property is set to 1, then the user interface does not use the RMI registry. The RMI registry is only needed if the RMI client program is used.

When the EAR file is deployed, the Java objects search for an RMI registry in the current memory area. If no RMI registry is found, then a registry is created and is bound to the product instance.

In a clustered environment, creation of an RMI registry that is bound to an instance of the product can be problematic. If the JVM fails, then the other JVMs on the same physical server are not reachable by client programs.

The solution in a clustered environment is to deploy the RMI registry file, `rmi-reg.war`, on the application server. The `rmi-reg.war` file is deployed on a separate server and creates the registry independent of any product JVMs. If a JVM is shut down or recycled, the RMI communication is not lost. In a clustered environment that spans multiple physical servers, RMI must be deployed one time on every server. But if the `mxs.allowLocalObjects` property is set to 1 and you do not use the RMI client program in your environment, you do not need to deploy the `rmi-reg.war` file.

Related tasks:

“Deploying the remote method invocation registry file for WebLogic Server” on page 26

When you build the remote method invocation (RMI) registry file, the file must be deployed each physical server. Deployment includes creating an RMI registry service and, creating a batch file to start RMI. To run the process correctly, you must update the start sequence for all servers so that all product servers start after the RMI server starts.

“Deploying the remote method invocation registry file in WebSphere Application Server” on page 14

A server that has a remote method invocation (RMI) registry continues to run even if another server in the cluster fails. When you build the RMI registry file, the file must be deployed in the application server.

“Building the RMI registry file” on page 14

The `rmi-reg.war` file is used to create the remote method invocation (RMI) registry. After you create the `rmi-reg.war` file, you can deploy the file on the application server.

Creating clusters in WebSphere Application Server

In the Integrated Solutions Console, you can create as many clusters as your deployment requires. Each cluster can consist of two or more Java virtual machine (JVM) cluster members.

Procedure

1. Log in as the administrative user to the Integrated Solutions Console by using the address `http://servername:9060/admin`
2. In the navigation pane, select **Servers > Clusters > WebSphere application server clusters**
3. Click **New** and enter a name for the cluster, such as `uiccluster`.
4. Specify the name of the first JVM cluster member, such as `maximoui1`.
5. Select **MXServer** to create the JVM cluster member based on an existing JVM.
6. Repeat steps 4-5 to create as many JVM cluster members as your environment requires.
7. Save your changes. The JVM cluster members are created as application servers, based on the settings in the MXServer application server.
8. Define the JVM parameter `-Dmxs.name` with the `jvmname` for each JVM cluster member.

- a. In the navigation pane, select **Servers > Server Types > WebSphere application servers** and select a JVM cluster member.
- b. Under **Service Infrastructure**, click **Java and Process Management > Process definition > Java Virtual Machine**.
- c. Scroll down and type 1536 for **Initial Heap Size** and 4096 for **Maximum Heap Size**.
- d. In the **Generic JVM arguments** field, specify the JVM parameter **Dmxe.name** to name each server, for example:
-Dmxe.name=maximoui1 **-Dmxe.name** is the property name passed to the JVM at startup time and **maximoui1** is the name that you identify as the JVM cluster member.
- e. Repeat steps a-d for each JVM cluster member.
9. In the navigation pane, select **Servers > Server Types > WebSphere application servers** and obtain the port numbers of each member:
 - a. Open the configuration of the JVM cluster member.
 - b. In the Communications section, click **Ports**.
 - c. Record the port number for the WC_defaulthost Port Name.
 - d. Repeat steps a-c for each JVM cluster member.
10. For each JVM cluster member, register the port number on a virtual host.
 - a. In the navigation pane, select **Environment > Virtual hosts**.
 - b. To create a virtual host, click **New** and specify the virtual host name.
 - c. Save your changes.
 - d. To create a host alias, select the virtual host and click **New**.
 - e. Enter the web server port number. Leave * as the host name.
 - f. Click **New** and enter the port of the cluster member. Repeat this step for each cluster member. Leave * as the host name.
 - g. Save your changes.
11. Repeat steps 2-10 for each cluster in your deployment.

Related concepts:

“Performance improvements with clusters” on page 2

A cluster groups similar functions on two or more Java virtual machines (JVMs) to process a single function, such as scheduled cron tasks. Clusters connect to the same database but operate independently. For example, if the cron task cluster fails, users can still connect to the user interface cluster.

Configuring Java Message Service for WebSphere Application Server

Java Message Service (JMS) is the messaging standard that is used to send and receive messages from queues. This process enables distributed communication with external systems in a loosely coupled, reliable, and asynchronous manner. The JMS configuration is application server-specific. You must configure JMS queues within the environment and make them accessible through the Java Naming Directory Interface (JNDI).

Before you begin

Configuration of JMS requires the understanding of buses, connection factories, queues, activation specifications, and stores. To set up the JMS configuration, you must be familiar with the configuration details for your application server.

About this task

Integration with external systems is supported through two message order processing mechanisms that use message queues. The first is sequential message processing, where the message order is guaranteed. The second is continuous message processing, where the messages are processed in parallel with message-driven beans (MDBs). The order in which messages are processed by this mechanism is not guaranteed.

If you are setting up an environment with an integration cluster that is connected to one or more external systems, you need to configure the JMS for queue-based integration. You need to create queues that are accessible by the user interface, cron, and integration clusters. The report cluster does not require JMS queues.

Related concepts:

“Java Message Service configuration for WebSphere Application Server” on page 23
If you are setting up an environment with an integration cluster that is connected to one or more external systems, you need to configure the Java Message Service (JMS) for queue-based integration. You need to create queues that are accessible by the user interface, cron, and integration clusters. The report cluster does not require JMS queues.

Creating data source providers and data sources:

Before you create service integration buses, you must create a data source provider. If you choose to use a database for the queue data, you must also create a data source.

About this task

Each data source requires a unique schema. Because of possible interference with scheduled backups, do not use the Maximo schema as a data source. When you add bus members, a unique message engine is created.

To ensure that the data source is available to all clusters that require access, create the data source provider at the cell level. To complete the configuration of the data source, you need information about your database configuration, such as path name and JDBC drive.

Procedure

1. In the database, create and configure that database that you plan to use as the data source.
2. In WebSphere Application Server, configure the J2C Authentication data and the JDBC provider for the data source
3. Test the connection to the data source.

Related concepts:

“Java Message Service configuration for WebSphere Application Server” on page 23
If you are setting up an environment with an integration cluster that is connected to one or more external systems, you need to configure the Java Message Service (JMS) for queue-based integration. You need to create queues that are accessible by the user interface, cron, and integration clusters. The report cluster does not require JMS queues.

Creating buses for Java Message Service:

A service integration bus consists of member application servers that share a common infrastructure to exchange information.

About this task

The naming convention for bus names combines the following identifiers:

- The function of the bus, such as mif for Maximo Integration Framework, ui for user interface, and cron for cron tasks.
- The service that uses the bus, such as jms for Java Message Service (JMS)
- The message engine type, such as bus for a service integration bus

Procedure

1. In the navigation pane of the Integrated Solutions Console, select **Service Integration > Buses**, and create a bus and specify uijmsbus for the name.
2. Add the user interface cluster as a member of the uijmsbus bus. By default, adding clusters to the bus creates a message engine for each cluster.
3. Create a bus and specify mifjmsbus for the name.
4. Add the integration framework cluster as a member of the mifjmsbus bus.
5. Optional: If you plan to send integration messages from cron tasks that create or update business objects in Maximo Asset Management, such as work order generation or reorder, create a bus and specify cronjmsbus for the name.
6. Optional: Add the cron task cluster as a member of the cronjmsbus bus.

Related concepts:

“Java Message Service configuration for WebSphere Application Server” on page 23
If you are setting up an environment with an integration cluster that is connected to one or more external systems, you need to configure the Java Message Service (JMS) for queue-based integration. You need to create queues that are accessible by the user interface, cron, and integration clusters. The report cluster does not require JMS queues.

Creating connection factories for Java Message Service:

Connection factories are used by the bus to create connections with Java Message Service (JMS) providers. Connection factories are stored in a Java Naming and Directory Interface (JNDI) namespace. You create a connection factory for each of the buses.

Before you begin

For each bus, look up the name of the message engine that was created automatically when the bus member was added to the bus. The default name format is *cluster_name.nnn-bus_name*. When you create connection factories, you must specify the name of the message engine in the **Target** field.

Procedure

1. Create a connection factory for the uijmsbus bus, and specify the following values:

Option	Description
Name	uiconfact
JNDI name	jms/maximo/int/cf/intcf
Bus name	uijmsbus

Option	Description
Target	The name of the message engine, which has the default format <i>cluster_name.nnn-bus_name</i>
Target type	Message Engine
Target significance	Required

2. Increase the maximum connections for the uiconfact connection factory from 10 to 50, depending on the expected load. The maximum number of connections depends on the load of transactions that you expect the JMS queues to process. If you expect a heavy transaction load, select a high number of maximum connections.
3. At the cell scope, create a connection factory for the mifjmsbus bus, and specify the following values:

Option	Description
Name	mifconfact
JNDI name	jms/maximo/int/cf/intcf
Bus name	mifjmsbus
Target	The name of the message engine, which has the default format <i>cluster_name.nnn-bus_name</i>
Target type	Message Engine
Target significance	Required

4. Increase the maximum connections for the mifconfact connection factory from 10 to 50, depending on the load.
5. Optional: If you created the cronjmsbus bus, create a connection factory for the cronjmsbus bus, and specify the following values:

Option	Description
Name	cronconfact
JNDI name	jms/maximo/int/cf/intcf
Bus name	cronjmsbus
Target	The name of the message engine, which has the default format <i>cluster_name.nnn-bus_name</i>
Target type	Message Engine
Target significance	Required

6. Optional: If you created a connection factory for the cronjmsbus bus, increase the maximum connections for the cronconfact connection factory from 10 to 50, depending on the load.

Related concepts:

"Java Message Service configuration for WebSphere Application Server" on page 23

If you are setting up an environment with an integration cluster that is connected to one or more external systems, you need to configure the Java Message Service (JMS) for queue-based integration. You need to create queues that are accessible by the user interface, cron, and integration clusters. The report cluster does not require JMS queues.

Creating queue destinations for Java Message Service:

You need to create queue bus destinations for each service integration bus. After the queue bus destinations are created, the queues are created for each destination.

Procedure

1. For the user interface bus, create a queue bus destination named `sqoutuibd`.
The bus destination is required to support processing of messages through the outbound sequential queue.
2. Optional: If you created the service integration `cronjmsbus` bus for the cron task cluster, create a queue bus destination named `sqoutcronbd`.
3. For the integration cluster bus member, which is the `mifjmsbus` bus, create multiple bus destinations. Multiple bus destinations are required to support processing of messages through the inbound and outbound queues.
 - a. To support processing of messages through the outbound sequential queue, create a queue bus destination named `sqoutmifbd`.
 - b. To support processing of messages through the inbound sequential queue, create a queue bus destination named `sqinmifbd`.
 - c. To support processing of messages through the inbound continuous queue, create a queue bus destination named `cqinmifbd`.
 - d. To support processing of messages through the inbound continuous error queue, create a queue bus destination named `cqinerrmifbd`.

Results

Based on the configuration settings, you can now perform data imports from the integration framework cluster, which has the message-driven beans enabled. You can also now perform data exports from the user interface cluster. If you require data import or data export in additional clusters, you can change the configuration for the other clusters.

Related concepts:

“Java Message Service configuration for WebSphere Application Server” on page 23
If you are setting up an environment with an integration cluster that is connected to one or more external systems, you need to configure the Java Message Service (JMS) for queue-based integration. You need to create queues that are accessible by the user interface, cron, and integration clusters. The report cluster does not require JMS queues.

Creating queues for Java Message Service:

You configure queues based on the queue destinations, which are used to send and receive messages to queues. A queue bus destination defines the bus name and queue name for the queue. You create one queue for each destination based on the default Java Message Service (JMS) provider.

Procedure

1. For the outbound sequential queue, create the queue for the user interface bus member with the following values:

Option	Description
Name	<code>sqoutui</code>
Bus name	<code>uijmsbus</code>
JNDI name	<code>jms/maximo/int/queues/sqout</code>
Queue name	<code>sqoutuibd</code>

2. Optional: If you created a cron task bus, then for the outbound sequential queue, create the queue for the cron task bus member with the following values:

Option	Description
Name	sqoutcron
Bus name	cronjmsbus
JNDI name	ms/maximo/int/queues/sqout
Queue name	sqoutcronbd

3. For the outbound sequential queue, create the queue for the integration framework bus member with the following values:

Option	Description
Name	sqoutmif
Bus name	mifjmsbus
JNDI name	jms/maximo/int/queues/sqout
Queue name	sqoutmifbd

4. For the inbound sequential queue, create the queue for the integration framework bus member with the following values:

Option	Description
Name	sqinmif
Bus name	mifjmsbus
JNDI name	jms/maximo/int/queues/sqin
Queue name	sqinmifbd

5. For the continuous queue inbound, create the queue for the integration framework bus member with the following values:

Option	Description
Name	cqinmif
Bus name	mifjmsbus
JNDI name	jms/maximo/int/queues/cqin
Queue name	cqinmifbd

6. For the inbound continuous error queue, create the queue for the integration framework bus member with the following values:

Option	Description
Name	cqinerrmif
Bus name	mifjmsbus
JNDI name	jms/maximo/int/queues/cqinerr
Queue name	cqinerrmifbd

Related concepts:

“Java Message Service configuration for WebSphere Application Server”

If you are setting up an environment with an integration cluster that is connected to one or more external systems, you need to configure the Java Message Service (JMS) for queue-based integration. You need to create queues that are accessible by the user interface, cron, and integration clusters. The report cluster does not require JMS queues.

Creating Java Message Service activation specifications:

A Java Message Service (JMS) activation specification is associated with a queue that uses message-driven beans (MDBs) to consume messages from the queue. The activation specification provides the information necessary for the queue to receive messages. For each continuous queue that you create, set up the activation specification at the cell scope.

Procedure

1. For the continuous queue inbound, create the activation specification with the following values:

Option	Description
Name	intjmsact
Bus name	mifjmsbus
Destination Type	queue
Destination JNDI Name	ms/maximo/int/queues/cqin

2. For the error queue, create the activation specification with the following values:

Option	Description
Name	intjmsacterr
Bus name	mifjmsbus
Destination Type	queue
Destination JNDI Name	jms/maximo/int/queues/cqinerr

What to do next

To complete the configuration of JMS, in the External Systems application, you must update the configuration of the JMS queues to reflect the JNDI names created for the connection factory and queues.

Related concepts:

“Java Message Service configuration for WebSphere Application Server”

If you are setting up an environment with an integration cluster that is connected to one or more external systems, you need to configure the Java Message Service (JMS) for queue-based integration. You need to create queues that are accessible by the user interface, cron, and integration clusters. The report cluster does not require JMS queues.

Java Message Service configuration for WebSphere Application Server:

If you are setting up an environment with an integration cluster that is connected to one or more external systems, you need to configure the Java Message Service

(JMS) for queue-based integration. You need to create queues that are accessible by the user interface, cron, and integration clusters. The report cluster does not require JMS queues.

Integration with external systems using message queues is supported through two default message order processing mechanisms. The first is sequential message processing, where the message order is guaranteed. The second is continuous message processing, where the messages are processed in parallel for better performance. The order in which messages are processed by this mechanism is not guaranteed.

When you use the continuous message processing, some messages that depend on a certain order can fail. For example, a vendor purchase order is processed before the vendor record is added. This processing order can prevent the purchase order from being processed. However, if the purchase order is reprocessed after the vendor record is added, the purchase order message is processed successfully.

The continuous message processing uses message-driven beans (MDBs) to process messages in a multi-threaded mode. There can be cases when the number of messages in error might reach a limit such that all MDBs continuously process only messages in error. This results in the number of messages in the queue to grow as no messages are processed successfully and removed from the queue. The limit for the number of error messages is equal to or greater than the maximum batch size of the queue multiplied by the number of MDBs deployed. In order to avoid this condition, configure the continuous queue with a corresponding error queue (exception destination). This configuration moves the messages in error to a different queue and allow new messages received into the queue to be processed.

The following table outlines the default setup for integration queues:

Table 1. Queues for a WebSphere Application Server setup

Queue	Description
Sequential inbound queue	Data comes in from external systems and is processed in the order in which the data is received.
Sequential outbound queue	Data goes out of the system to external systems in the order in which the data is processed by the system.
Continuous inbound queue	Data comes into the system from external systems that does not need to be processed in the order that the data is received. Messages can be processed in parallel by multiple MDBs.
Continuous inbound error queue	Error messages that result from the continuous inbound queue are placed in this queue for message reprocessing and error handling.

Related concepts:

Access to services by inbound messages

Related tasks:

“Configuring Java Message Service for WebSphere Application Server” on page 17
Java Message Service (JMS) is the messaging standard that is used to send and receive messages from queues. This process enables distributed communication with external systems in a loosely coupled, reliable, and asynchronous manner. The JMS configuration is application server-specific. You must configure JMS queues within the environment and make them accessible through the Java Naming Directory Interface (JNDI).

Configuring a message processing server

Deploying EAR files for clusters in WebSphere Application Server

When you deploy the EAR files in WebSphere Application Server, the changes that you made for the different clusters, such as the configuration of the message-driven beans, are deployed.

About this task

The *install_home* variable represents the installed location of the Maximo Asset Management folder, which by default is `ibm\SMP\maximo`.

Procedure

1. Stop any Java virtual machines (JVM) that are running in the cluster in which you want to deploy EAR files.
2. Open the Integrated Solutions Console and in the navigation pane, click **Applications > Application Types > WebSphere enterprise applications**.
3. Click **Install** and in the *install_home/deployment/default* directory, locate the EAR file that you want to deploy. For example, if you are setting up the user interface cluster, locate the **maximoui.ear** file.
4. Accept the default settings and then select all the modules and the cluster, such as UICluster. If you are using a web server, select the web server.
5. Select all modules and select the virtual host, such as UICluster_host, to map the virtual hosts.
6. Click **Finish**.
7. Save the file to the master configuration.
8. Deploy the `maximoiehs.ear` file.
9. Repeat steps 2-8 for the EAR files for the remaining clusters.

What to do next

Start the JVM for the remote method invocation registry, and then start the clusters in the application server. Log in to verify that the process is successful.

Related tasks:

“Building and deploying EAR files for basic configurations” on page 36
You can build and deploy EAR files for a basic configuration. In a clustered configuration, each cluster has its own EAR file to build and deploy.

Creating and deploying clusters in WebLogic Server

In WebLogic Server, you deploy the remote method invocation (RMI), create the clusters, configure Java Message Service (JMS), and deploy the EAR files.

Related concepts:

“Application server documentation” on page 43

For more information about your application server, see the following web sites.

Deploying the remote method invocation registry file for WebLogic Server

When you build the remote method invocation (RMI) registry file, the file must be deployed each physical server. Deployment includes creating an RMI registry service and, creating a batch file to start RMI. To run the process correctly, you must update the start sequence for all servers so that all product servers start after the RMI server starts.

Related concepts:

“Java remote method invocation” on page 15

Remote method invocation (RMI) is an application programming interface that provides a way for objects in separate memory areas to interact. Separate memory areas can be part of the same physical system or can be on different systems connected by a network.

Java remote method invocation:

Remote method invocation (RMI) is an application programming interface that provides a way for objects in separate memory areas to interact. Separate memory areas can be part of the same physical system or can be on different systems connected by a network.

An RMI Registry is an area in memory that maintains the RMI address information of a Java object server. By default, the RMI registry is created on port 1099. More than one RMI registry can exist in memory. Each registry has a designated TCP/IP port for access.

If the **mxe.allowLocalObjects** property is set to 1, then the user interface does not use the RMI registry. The RMI registry is only needed if the RMI client program is used.

When the EAR file is deployed, the Java objects search for an RMI registry in the current memory area. If no RMI registry is found, then a registry is created and is bound to the product instance.

In a clustered environment, creation of an RMI registry that is bound to an instance of the product can be problematic. If the JVM fails, then the other JVMs on the same physical server are not reachable by client programs.

The solution in a clustered environment is to deploy the RMI registry file, **rmi.reg.war**, on the application server. The **rmi.reg.war** file is deployed on a separate server and creates the registry independent of any product JVMs. If a JVM is shut down or recycled, the RMI communication is not lost. In a clustered environment that spans multiple physical servers, RMI must be deployed one time on every server. But if the **mxe.allowLocalObjects** property is set to 1 and you do not use the RMI client program in your environment, you do not need to deploy the **rmi.reg.war** file.

Related tasks:

“Deploying the remote method invocation registry file for WebLogic Server” on page 26

When you build the remote method invocation (RMI) registry file, the file must be deployed each physical server. Deployment includes creating an RMI registry service and, creating a batch file to start RMI. To run the process correctly, you must update the start sequence for all servers so that all product servers start after the RMI server starts.

“Deploying the remote method invocation registry file in WebSphere Application Server” on page 14

A server that has a remote method invocation (RMI) registry continues to run even if another server in the cluster fails. When you build the RMI registry file, the file must be deployed in the application server.

“Building the RMI registry file” on page 14

The `rmireg.war` file is used to create the remote method invocation (RMI) registry. After you create the `rmireg.war` file, you can deploy the file on the application server.

Creating remote method invocation registry services for WebLogic Server:

You can deploy a remote method invocation (RMI) registry file to create a registry that is independent of the product servers. A server that contains an RMI registry continues to run even if another server in the cluster fails. This registry starts on the server before any of the cluster members starts.

About this task

The `install_home` variable represents the installed location of the product folder. The `WebLogic_install_home` variable represents the installed location of WebLogic Server, which by default is `\bea`.

Procedure

1. Open a command prompt and navigate to the `WebLogic_install_home\user_projects\domains\domain_name` directory.
2. Run the `startWebLogic.cmd` file to start the WebLogic Server.
3. Open the WebLogic Server administrative console. The default URL is `http://servername:7001/console`.
4. In the navigation pane, browse to the **domain_name > Servers** folder.
5. Click **Configure New Server** and in the **Name** field, specify `RMIRegistry`. Spaces are invalid characters.
6. In the **Listen Port** field, specify `9999` and click **Create**.
7. In the navigation pane, click **Deployments > Web Application Modules > Deploy a New Web Application Module**.
8. Specify `install_home/deployment/default` as your archive directory.
9. Select the `rmireg.war` file and use the default name, or specify a different name for the file.
10. Click **Finish** to deploy the file.
11. Click **Save** and close the console.

What to do next

Create the batch file that starts the RMI registry.

Creating a batch file to start remote method invocation on WebLogic Server:

After you create the remote method invocation (RMI) registry service, you must create a batch file to start RMI on the application server.

Procedure

1. Navigate to the *WebLogic_install_home/user_projects/domains/domain_name* directory.
2. Create a backup of the *startWebLogic.cmd* file, then rename the *startWebLogic.cmd* file to *startRMIRegistry.cmd*.
3. Edit the *startRMIRegistry.cmd* file and change the **SERVER_NAME** parameter to *RMIRegistry*.
4. Add or modify the *set MEM_ARGS* code to match the following code: *set MEM_ARGS=-Xms5m -Xmx10m*.
5. Save and close the file.
6. Open a new command prompt and navigate to the *WebLogic_install_home/user_projects/domains/domain_name* directory.

Creating clusters in WebLogic Server

You can create as many clusters as your deployment requires.

Procedure

1. Start the application server.
 - a. In a command prompt, change to the *bea\user_projects\domains\base_domain* directory.
 - b. Run the **startweblogic** command.
2. Log in as the administrative user to the WebLogic Server administration console with the address *http://servername:7001/console*.
3. To edit within the administration console, lock the configuration edit hierarchy for the domain.
4. Create the managed servers that you plan to add to the cluster.
5. Create the cluster.
6. Select the cluster, and, in the **Servers** tab, add the managed servers to the cluster.
7. Set the minimum heap size to 128 MB.
8. Set the maximum heap size to 1424 MB.
9. Assign a port to the cluster. Each cluster must have a unique port.
10. Repeat steps 4-9 for each cluster that your deployment requires.
11. Activate your changes.

Related concepts:

“Performance improvements with clusters” on page 2

A cluster groups similar functions on two or more Java virtual machines (JVMs) to process a single function, such as scheduled cron tasks. Clusters connect to the same database but operate independently. For example, if the cron task cluster fails, users can still connect to the user interface cluster.

Configuring the Java Message Service for WebLogic Server

Java Message Service (JMS) is used as the messaging standard to create, send, receive, and read messages from queues. This process enables distributed communication with external systems in a loosely coupled, reliable, and asynchronous manner.

Before you begin

Configuration of JMS requires the understanding of buses, connection factories, queues, activation specifications, and stores. To set up the JMS configuration, you must be familiar with the configuration details for your application server.

Related concepts:

“Java Message Service configuration for WebLogic Server” on page 34

You use Java Message Service (JMS) servers to manage queue and topic resources, and to maintain information about queue stores.

Creating queue stores:

You create Java Message Service (JMS) stores to store the persistent messages. For sequential outbound queues, you create the queue stores for each cluster to ensure that specific cluster downtime does not affect other clusters. For other queue types (sequential inbound queue and continuous inbound queue), you ensure that they are configured for the integration framework and cron task clusters.

Procedure

1. Create a sequential outbound queue store, and name the store `sqoutuistor`. Target the `sqoutuistor` store to one of the user interface cluster servers.
2. Create the following stores for each queue that you are creating. Target each store to one of the integration framework cluster servers.
 - a. For the sequential inbound queue, create a store and name it `sqinstore`.
 - b. For the sequential outbound queue, create a store and name it `sqoutintstore`.
 - c. For the continuous inbound queue, create a store and name it `cqinstore`.
3. Create a sequential outbound queue store, and name the store `sqoutcronstore`. Target the `sqoutcronstore` store to one of the cron task cluster servers.

Related concepts:

“Java Message Service configuration for WebLogic Server” on page 34

You use Java Message Service (JMS) servers to manage queue and topic resources, and to maintain information about queue stores.

Creating Java Message Service servers:

Java Message Service (JMS) servers manage JMS queue and topic resources. The queue and topic resources are defined within JMS modules that are targeted to a specific JMS server. JMS servers also maintain information about the store that you use for persistent messages that are received on queue destinations.

Procedure

1. For the sequential outbound queue, sqoutuistore, create a JMS server and name it sqoutuiserver. Target the server to the user interface cluster server.
2. For the integration framework cluster, create the following JMS servers. Target the servers to the integration framework cluster server.
 - a. For the sequential inbound queue, create a JMS server for the sqinstore store, and name the queue sqinserver.
 - b. For the sequential outbound queue, create a JMS server for the sqoutinstore store, and name the queue sqoutintserver.
 - c. For the continuous inbound queue, create a JMS server for the cqinstore store, and name the server cqinserver.

For this server, set the maximum bytes to a value based on your JVM maximum heap size. This value is typically set to approximately 10% to 20% of the maximum heap size, which prevents memory errors if messages are created faster than the consumer can process.
3. For the sequential outbound queue, create a JMS server for the sqoutcronstore store and name the server sqoutcronserver. Target the sqoutcronserver server to one of the cron task cluster servers.

Related concepts:

“Java Message Service configuration for WebLogic Server” on page 34

You use Java Message Service (JMS) servers to manage queue and topic resources, and to maintain information about queue stores.

Creating Java Message Service modules:

Java Message Service (JMS) modules are configuration containers for JMS resources. The JMS modules store the information for the connection factories that queues are configured to use.

Procedure

1. Create a JMS module for the sequential outbound queue, and name the module intjmssqoutuimodule. Target the intjmssqoutuimodule module to the user interface cluster.
 - a. Create a sequential outbound queue for the intjmssqoutuimodule module with the following values:

Name sqout

JNDI Name
jms/maximo/int/queues/sqout

Using the default name, sqout, create a subdeployment for the sequential outbound queue that you created, and target it to the sqoutuiserver JMS server.
 - b. Create a connection factory for the queues that you created with the following values:

Name intjmssqconfact

JNDI Name
jms/maximo/int/cf/intsqcf

Do not create a subdeployment because the connection factory inherits the JMS module target, which is the integration framework cluster.

Set the connection factory XA transaction to enabled.

2. Create a JMS module for the sequential queues, and name the module `intjmssqintmodule`. Target the `intjmssqintmodule` module to the integration framework cluster.
 - a. Create a sequential inbound queue for the `intjmssqintmodule` module with the following values:

Name `sqin`

JNDI Name
`jms/maximo/int/queues/sqin`

Using the default name, `sqin`, create a subdeployment for the sequential inbound queue that you created, and target it to the `sqinserver` JMS server.
 - b. Create a sequential outbound queue for the `intjmssqintmodule` module with the following values:

Name `sqout`

JNDI Name
`jms/maximo/int/queues/sqout`

Using the default name, `sqout`, create a subdeployment for the sequential outbound queue that you created, and target it to the `sqoutserver` JMS server.
 - c. Create a connection factory for the queues that you created with the following values:

Name `intjmssqconfact`

JNDI Name
`jms/maximo/int/cf/intsqcf`

Do not create a subdeployment because the connection factory inherits the JMS module target, which is the integration framework cluster.

Set the connection factory XA transaction to enabled.
3. Create a JMS module for the sequential outbound queue, and name the module `intjmssqoutcronmodule`. Target the `intjmssqoutcronmodule` module to the cron task cluster.
 - a. Create a sequential outbound queue for the `intjmssqintmodule` module, with the following values:

Name `sqout`

JNDI Name
`jms/maximo/int/queues/sqout`

Using the default name, `sqout`, create a subdeployment for the sequential outbound queue, and target it to the `sqoutcronserver` JMS server.
 - b. Create a connection factory for the queues that you created with the following values:

Name `intjmssqconfact`

JNDI Name
`jms/maximo/int/cf/intsqcf`

Do not create a subdeployment because the connection factory inherits the JMS module target, which is the integration framework cluster. Set the connection factory XA transaction to enabled.
4. Create a JMS module for the continuous queue and name it `intjmsscmodule`. Target the `intjmsscmodule` module to the integration framework cluster.

- a. Create a continuous inbound queue for intjmscqmodule with the following values:

Name cqin

JNDI Name

jms/maximo/int/queues/cqin

Create a subdeployment for this queue with the default name, cqin, and target it to the cqinserver JMS server.

- b. Create a connection factory for the queue that you created with the following values:

Name intjmscqconfact

JNDI Name

jms/maximo/int/cf/intcqcf

Do not create a subdeployment; the connection factory inherits the JMS module target, which is the integration framework cluster.

Set the connection factory XA transaction to enabled. Set the **Messages Maximum** field to -1.

Related concepts:

“Java Message Service configuration for WebLogic Server” on page 34

You use Java Message Service (JMS) servers to manage queue and topic resources, and to maintain information about queue stores.

Creating Java Database Connectivity data sources in WebLogic Server:

When you create a data source provider, you must create the Java Database Connectivity (JDBC) data source, which specifies the connection information for the database.

Procedure

1. Log into the WebLogic Server administration console <http://servername:7001/console> and click **Lock and Edit**.
2. Select **Services > JDBC > Data Sources**.
3. Create a data source. Specify jmsqueuedatasource as the name, and specify the JNDI name.
4. Specify the database type and the database driver. This driver must be a non-XA driver.
5. Clear the **Supports Global Transactions** check box.
6. Specify the database name, the host name, the port number, and the user information.
7. Test the connection.
8. Update the information if needed.
9. Select all the servers needed for the cluster and click **Finish**.
10. Activate your changes.

Related concepts:

“Java Message Service configuration for WebLogic Server” on page 34

You use Java Message Service (JMS) servers to manage queue and topic resources, and to maintain information about queue stores.

Related tasks:

Manually creating a data source for the persistent store

Creating data stores for WebLogic Server:

When you create data sources, you must create a data store for each of the four queues. Stores are used to hold the queue messages.

About this task

The prefix values are important. If you do not have unique names for prefix values on each store, messages can become corrupted because they all use the same store files.

Procedure

1. In the WebLogic Server administrative console, click **Lock and Edit**.
2. Select **Services > Persistent stores**.
3. Create a Java Database Connectivity (JDBC) store for a continuous inbound queue.
4. Specify the name as **mxintcqinstore**.
5. Specify the target as **cqinserver**.
6. Select the data source you created previously and enter the prefix value **mxintcqin**.
7. Click **Finish**.
8. Repeat steps 4-7 to create a JMS JDBC store for a sequential inbound queue. Specify the following values:

Option	Description
Name	mxintsqinstore
Target	MAXIMOIF
Prefix value	mxintsqin

9. Repeat steps 4-7 to create a JMS JDBC store for a sequential outbound queue. Specify the following values:

Option	Description
Name	mxintsqoutstore
Target	MAXIMOU11
Prefix value	mxintsqout

10. Repeat steps 4-7 to create a JMS JDBC store for a continuous inbound error queue. Specify the following values:

Option	Description
Name	mxintcqinerrstore
Target	MAXIMOIF

Option	Description
Prefix value	mxintcginerr

11. Activate your changes.

Related concepts:

“Java Message Service configuration for WebLogic Server”

You use Java Message Service (JMS) servers to manage queue and topic resources, and to maintain information about queue stores.

Creating Java Message Service connection factories in WebLogic Server:

When you create a Java Message Service (JMS) module, you must create a connection factory to access bus destinations. The connection factory specifies the Java Naming and Directory Interface (JNDI) name.

Procedure

1. In the WebLogic Server administration console <http://servername:7001/console>, click **Lock and Edit**.
2. Select **Services > Messaging > JMS Modules**.
3. Select the JMS module to which you want to create a connection factory and click **New**.
4. Click **Connection Factory** and click **Next**.
5. Specify `intjmsconfact` as the name.
6. Specify `jms/maximo/int/cf/intcf` as the JNDI name.
7. Accept the default settings for the targets and click **Finish**.

Related concepts:

“Java Message Service configuration for WebLogic Server”

You use Java Message Service (JMS) servers to manage queue and topic resources, and to maintain information about queue stores.

Activating Java Message Service connection factories in WebLogic Server:

When you create connection factories, you must activate them.

Procedure

1. In the WebLogic Server administration console, click **Lock and Edit**.
2. Select **Services > Messaging > JMS Modules**.
3. Select the connection factory that you want to activate.
4. On the **Transaction** tab, select the **XA Connection Factory Enabled** check box.
5. Click **Save**.
6. On the **Client** tab, set the maximum number of messages per session to -1.
7. Save and activate the changes.

Related concepts:

“Java Message Service configuration for WebLogic Server”

You use Java Message Service (JMS) servers to manage queue and topic resources, and to maintain information about queue stores.

Java Message Service configuration for WebLogic Server:

You use Java Message Service (JMS) servers to manage queue and topic resources, and to maintain information about queue stores.

JMS servers and modules provide messaging support. If you use WebLogic Server, you cannot use the loop-back messaging technique for errors that occur in the continuous message processing mode. With WebLogic Server, you cannot set an error queue to be the error queue for itself. Therefore, the error queue clogs after a few errors, depending on the maximum messages per session value of the connection factory.

Because the WebLogic Server error queue clogs after a few errors, only the errors in the front of the queue are processed. The remaining messages are not processed, unless the error messages are deleted. To avoid this issue, set the maximum messages per session value to -1 for the continuous queue connection factory. The -1 value indicates that there is no limit on the number of messages. However, the number of messages is still limited by the amount of remaining virtual storage for the process.

You can set up the following queues:

Table 2. Queues for a WebLogic Server setup

Queue	Description
Sequential inbound queue	Data comes into the system from external systems that must be run in the order that it is received
Sequential outbound queue	Data goes out of the system to external systems
Continuous inbound queue	Data comes into the system from external systems that must not be processed in the order that it is received. It can be run in parallel by multiple message-driven beans (MDBs).

Related concepts:

Access to services by inbound messages

Related tasks:

“Configuring the Java Message Service for WebLogic Server” on page 29
Java Message Service (JMS) is used as the messaging standard to create, send, receive, and read messages from queues. This process enables distributed communication with external systems in a loosely coupled, reliable, and asynchronous manner.

Configuring a message processing server

Deploying EAR files for clusters in WebLogic Server

When you deploy the EAR files in WebLogic Server, the changes that you made for the different clusters, such as the configuration of message-driven beans, are deployed.

Procedure

1. To edit within the administrative console, lock the configuration edit hierarchy for the domain.
2. Select **Deployments**.

3. Select **Install** and browse to the location of the `maximo.ear` file, select the location, and save the file. The EAR file is saved in the `install_home\deployment\default` folder.
4. Select **Install** to deploy the EAR file to the cluster.
5. Activate the changes.
6. Deploy the `maximo.ear` file.
7. Repeat steps 2-6 for the remaining clusters.

What to do next

Start the JVM for the RMI registry, and then start the clusters in the application server. Log in to verify that the process is successful.

Related tasks:

“Building and deploying EAR files for basic configurations”

You can build and deploy EAR files for a basic configuration. In a clustered configuration, each cluster has its own EAR file to build and deploy.

Building and deploying EAR files for basic configurations

You can build and deploy EAR files for a basic configuration. In a clustered configuration, each cluster has its own EAR file to build and deploy.

Related tasks:

“Building Maximo EAR files for clusters” on page 11

After you create a build file for each cluster, you must build a Maximo EAR file for the cluster. The name of the EAR file is based on the set `EAR_FILENAME` statement in the build file.

“Deploying EAR files for clusters in WebLogic Server” on page 35

When you deploy the EAR files in WebLogic Server, the changes that you made for the different clusters, such as the configuration of message-driven beans, are deployed.

“Deploying EAR files for clusters in WebSphere Application Server” on page 25

When you deploy the EAR files in WebSphere Application Server, the changes that you made for the different clusters, such as the configuration of the message-driven beans, are deployed.

Building EAR files for basic configurations

The EAR files contains all the fields required to run an application. There are two EAR files: `maximo.ear` and `maximoiehs.ear`. Both the EAR files contain one or more web application modules. In a clustered configuration, each cluster requires a separate build file to create an EAR file for the cluster.

About this task

You can also build EAR files for example, to modify a database connection parameter in the `maximo.properties` file after the initial installation.

Procedure

1. Open a command prompt.
2. Go to `install_home\deployment`.
3. Run the appropriate script:

Option	Description
buildmaximoear (Windows)	Creates a maximo.ear file
buildmxiehsear	Creates a maximoiehs.ear file

These scripts take several minutes to run. The command prompt or terminal window displays a BUILD SUCCESSFUL line.

Related concepts:

“EAR files” on page 12

EAR files are archives that contain all the required files to run an application.

“Web application archive files” on page 13

Web application archive (WAR) files are part of EAR files. They contain, for example, JSP or HTML pages.

Related tasks:

“Building Maximo EAR files for clusters” on page 11

After you create a build file for each cluster, you must build a Maximo EAR file for the cluster. The name of the EAR file is based on the set *EAR_FILENAME* statement in the build file.

Deploying Maximo Enterprise Application Archive files in WebSphere Application Server

You can deploy EAR files in WebSphere Application Server.

Procedure

1. Open the Integrated Solutions Console and in the navigation pane, click **Applications > Application Types > WebSphere enterprise applications**.
2. Optional: If the Maximo and Maximoiehs applications are listed, select them and click **Uninstall** and click **OK**.
3. After the applications are uninstalled, return to the **WebSphere enterprise applications** screen and click **Install**.
4. Browse to the location of the maximo.ear file, select the location, and click **Next**. The Maximo EAR files are saved in the *install_home\deployment\default* folder.
5. Click **Next**.
6. In the Clusters and Servers section of the Map modules to servers window, select all components and click **Apply** and then **Next**.
7. In the Virtual host column of the Map virtual hosts for web modules window, select Maximo_host in the drop-down list for all web modules, and click **Next**.
8. Click **Finish**.
9. When the EAR file is deployed, save the file to the master configuration.

What to do next

After you deploy the Maximo.ear file, repeat the process to deploy the Maximoiehs.ear file. You must start the application server and login to verify that the installation has been successful.

Related concepts:

“EAR files” on page 12

EAR files are archives that contain all the required files to run an application.

Deploying Maximo Enterprise Application Archive files in WebLogic Server

You can deploy EAR files in a WebLogic Server.

Procedure

1. Open the WebSphere Administration Console and click **Lock & Edit**.
2. Click **Deployments**.
3. Optional: If the Maximo and Maximoiehs applications are listed, select them and click **Delete**.
4. After the applications are deleted, click **Install**.
5. Browse to the location of the maximo.ear file, select the location, and click **Next**. The Maximo ear files are saved in the *install_home*\deployment\default folder.
6. Click **Install** this deployment as an application and click **Next**.
7. Click **Finish**.
8. Activate the changes.

What to do next

After you deploy the Maximo.ear file, repeat the process to deploy the Maximoiehs.ear file. You must start the application server and login to verify that the installation has been successful.

Related concepts:

“EAR files” on page 12

EAR files are archives that contain all the required files to run an application.

Configuring general settings

Whether you configure a basic system or a clustered system, you can create Java virtual machines, configure Internet Explorer settings, and configure session timeout periods. You can also migrate the administrative workstation.

Online help configuration

There are different deployment options for online help running in an information center.

- You can deploy the online help application archive file (maximoiehs.ear) on the same server or cluster of servers on which the EAR file (maximo.ear) file is deployed.
- You can deploy the maximoiehs.ear file on a separate server. All application archive (maximo.ear) file deployments can see the help that runs on the separate server.

Regardless of how you deploy the information center, a set of system properties connects the product user interface (UI) to the information center. The values in these “mxe.help” system properties must match the deployed information center to make online help available from the UI.

Related reference:

“mxe.help properties” on page 433

The **mxe.help** system properties connect the user interface to the information center. Some of the properties are used to construct the link that opens the information center. To ensure that information center is available, match the values in the **mxe.help** properties to the information center that you deploy.

Web application archive files

Web application archive (WAR) files are part of EAR files. They contain, for example, JSP or HTML pages.

WAR file	Description
maximouiweb.war	Contains the user interface-related JavaServer Pages (.jsp files), Java classes, static HTML files, and static image files. The buildmaximoeear.xml file has information about the files in this module. This web application uses the configuration details in the web.xml file, located in the <maximo root>\applications\Maximo\Maximouiweb\webmodule\WEB-INF folder. This file also specifies the URL to access online help.
mboweb.war	Contains the business objects, Java classes, and dependent third-party Java classes.
meaweb.war	The integration framework enables the exchange of application data with another application or with an external system. Users can create and maintain data in one system and use the integration framework to transfer data to an external system, which eliminates duplicate processing.
iehs.war	Provides the online help. The buildmxiehsear.xml file has information about all the files in this module.
rmireg.war	Creates the remote method invocation (RMI) registry file.

Related concepts:

“EAR files” on page 12

EAR files are archives that contain all the required files to run an application.

Related tasks:

“Creating build files for clusters” on page 10

You must create a separate buildmaximoeear.cmd file for each cluster. When you run the separate buildmaximoeear.cmd files, you create a separate EAR file for each cluster.

“Building Maximo EAR files for clusters” on page 11

After you create a build file for each cluster, you must build a Maximo EAR file for the cluster. The name of the EAR file is based on the set *EAR_FILENAME* statement in the build file.

EAR files

EAR files are archives that contain all the required files to run an application.

The following two EAR files are used. Each EAR file contains one or more web application modules (.war extension):

- `maximo.ear`
 - `maximouiweb.war`
 - `mboweb.war`
 - `meaweb.war`
- `maximoiehs.ear`
 - `iehs.war`

You rebuild and redeploy EAR files whenever you:

- Modify .xml files or custom class files (`maximo.ear`).
- Modify HTML help topics (online help) (`maximoiehs.ear`).
- Modify settings in the `maximo.properties` file (`Maximo.ear`).

Related concepts:

“Web application archive files” on page 13

Web application archive (WAR) files are part of EAR files. They contain, for example, JSP or HTML pages.

Related tasks:

“Building EAR files for basic configurations” on page 36

The EAR files contains all the fields required to run an application. There are two EAR files: `maximo.ear` and `maximoiehs.ear`. Both the EAR files contain one or more web application modules. In a clustered configuration, each cluster requires a separate build file to create an EAR file for the cluster.

“Creating build files for clusters” on page 10

You must create a separate `buildmaximoeear.cmd` file for each cluster. When you run the separate `buildmaximoeear.cmd` files, you create a separate EAR file for each cluster.

“Building Maximo EAR files for clusters” on page 11

After you create a build file for each cluster, you must build a Maximo EAR file for the cluster. The name of the EAR file is based on the set `EAR_FILENAME` statement in the build file.

Configuring application servers

You can configure the memory settings for the application servers. You can also set up load balancing, enable secure socket layer support, and create Java virtual machines.

Memory settings for the application server process

The application server process in which the system is deployed must be configured with the right amount of memory setting, or else the process runs out of memory when the system is running.

A single process running the system can support up to 50 user loads with optimal performance. Scheduled cron jobs and integration activities within a process also consumes additional memory. A higher user load on a single process also can result in memory errors and can potentially cause the process to terminate.

The following recommended memory settings are for a single process that is running the system with a small amount of capacity for reporting, cron tasks, and integration activity. The same settings also apply to the application server processes that are set up to process the integration load or the cron tasks as part of a clustered configuration. An application server process can run into a memory

situation because of a large user load, large integration messages being processed, cron tasks that run for a long time and require more memory, bugs in the application code or the application server, and so on. When a memory situation occurs, identify the root cause. If the problem occurs because of a higher user load, adding additional servers helps.

WebLogic Server

If WebLogic Server is set up to run with JVM, use the following memory settings:

- Minimum heap size - 128 MB (-Xms512m)
- Maximum heap size - 1424 MB (-Xmx1424m)
- Maximum permanent size - 512 MB (-XX:MaxPermSize=512m)

WebSphere Application Server

If you are using a WebSphere Application Server, use the following memory settings for a 32 BIT JVM:

- Minimum heap size - 1536 MB (-Xms1424m)
- Maximum heap size - 1536 MB (-Xmx1424m)

If you are using a WebSphere Application Server, use the following memory settings for a 64 BIT JVM:

- Minimum heap size - 4096 MB (-Xms1424m)
- Maximum heap size - 4096 MB (-Xmx1424m)

Related concepts:

“Application server documentation” on page 43

For more information about your application server, see the following web sites.

Load balancing

Load balancing is the distribution of the task load across multiple instances of an application. A basic system configuration typically supports a user load of 50 users or less. A clustered configuration can support a larger user load.

User load comes from users who are logged in. Nonuser load comes from scheduled jobs (cron tasks) and incoming transactions from the integration framework. It is optional to distribute user load and nonuser load to different application servers or clusters.

For HTTP traffic, such as system applications, integration post, and so on, software load balancers and hardware load balancers are available. Typically, your application server vendor provides a load balancer option. A hardware load balancer generally provides better performance, but it is an additional expense. See the documentation specific to your application server for additional information.

Secure socket layer support

The system supports secure socket layout (SSL). For more information about how to enable SSL connectivity, see documentation specific to your application server.

Creating Java virtual machines

You can create a Java virtual machine (JVM) as part of a clustered configuration or to provide additional resources in a basic configuration. You create JVMs on the application server.

Creating Java virtual machines for WebSphere Application Server:

When you create a Java virtual machine (JVM), you can set the parameter and memory settings.

Procedure

1. Open the Integrated Solutions Console <http://servername:9060/admin> and log in.
2. In the navigation pane, click **Servers > New server**.
3. Click **WebSphere application server** and click **Next**.
4. Specify the server name and click **Next**.
5. Accept the default values for the server template and click **Next**.
6. Accept the default values for the server properties and click **Next**.
7. Click **Finish**, click **Save**, and then click **OK**.
8. Edit JVM memory settings and parameters:
 - a. Click **Servers** and then click the server you created.
 - b. Under **Service Infrastructure**, click **Java and Process Management > Process definition > Java Virtual Machine**.
 - c. Scroll down and type 1536 for **Initial Heap Size** and 4096 for **Maximum Heap Size**.
 - d. In the **Generic JVM arguments** field, use the JVM parameter `-Dmx.name=hostname~jvmname` to name each server. `-Dmx.name` is the property name passed to the JVM at startup time and `hostname~jvmname` is the name that you identify as the JVM server. With this information, when you look at the log file, a **donotrun** parameter, or a `maxsession` entry, you can identify the JVM.
For example, `-Dmx.name=computer1~uiserver1`
9. Set the new application server to start in running mode:
 - a. Click **Servers** and then click the server you created.
 - b. Click **Java and Process Management > Monitoring policy**.
 - c. Change the **Node restart state** to **RUNNING**.
 - d. Click **Apply** and then click **Save**.

Creating a Java virtual machine in WebLogic Server:

When you create a Java virtual machine (JVM), you can set the parameter and memory settings.

Procedure

1. In the WebLogic Server administration console <http://servername:7001/console>, click **Lock and Edit**.
2. In the navigation pane, click **Services > Messaging > JMS Modules**.
3. Click **Next**.
4. Specify the server name, the descriptor file name, and the location where the descriptor is stored, and click **Next**.
5. Specify the target server and cluster and click **Next**.
6. Specify whether you want to add resources to the JMS system module and click **Next**.
7. Click **Activate Changes**.

What to do next

Create resources for the JMS system module.

Application server documentation

For more information about your application server, see the following web sites.

WebSphere Application Server

For more information about WebSphere Application Server, see the IBM® WebSphere Application Server, Version 6.1 Information Center:

<http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/index.jsp>

WebLogic Server

For more information about WebLogic Server, see the Oracle WebLogic Server documentation:

<http://e-docs.bea.com/wls/docs92/> or WebLogic Server

Configuring browser settings

You can configure browser settings to ensure that the client browser checks for the current version of the page. You can also set session timeout periods for client browsers.

Configuring Internet Explorer settings

You must verify that the client browser checks for the current version of the page. You check this option through your Internet Explorer settings.

Procedure

1. From your web browser, select **Tools** and then **Internet Options**.
2. On the **General** tab, click **Settings**.
3. Select **Automatically**.
4. Click **OK**.

Configuring session timeout periods

By default, client sessions are timed out after 30 minutes of inactivity. To change this value, you can edit the web.xml file. Increasing the session-timeout element to a higher value consumes additional memory. It is recommended not to increase to a high value.

Procedure

1. Go to <Maximo_root>root>\applications\maximo\maximouiweb\webmodule\WEB-INF\web.xml.
2. Find the session-config section and change the session-timeout element to a different value. For example, replacing 30 with 60 increases the timeout period from 30 minutes to 60 minutes.

Migrating the administrative workstation

The administrative workstation can be migrated from one physical system to another.

About this task

In some cases, you might want to migrate an existing administrative workstation from one system to another. For example, an existing administrative system might be reassigned. You can delegate the deployment maintenance duties to another system.

The target system must host the same operating system and major version as the original administrative system. The target system must be of the same hardware type as the existing administrative workstation. The user ID used to install the product on the existing administrative workstation must also exist on the target system with the same permissions defined.

Procedure

1. Log on to the existing administrative system with the user ID used to install the product.
2. Create a copy of the installation directory. By default, this value is C:\IBM\SMP for Windows and /opt/IBM/SMP for Linux and UNIX systems. Ensure that all file permissions are preserved.
3. Log on to the target administrative system with the same user ID that was used to install the product on the existing administrative workstation.
4. Copy the installation files and directories to the file system of the target administrative system. You must maintain the directory structure of the original installation. For example, if the installation directory on the existing administrative system is C:\IBM\SMP, you cannot copy those files to a C:\NewAdminWS\IBM\SMP directory on the target administrative workstation.
5. Update the deployment engine host name by running the following commands:

Windows

```
install_directory\SMP\CTG_DE\acsi\bin\de_chghostname.cmd
```

Linux and UNIX

```
install_directory/SMP/CTG_DE/acsi/bin/de_chghostname.sh
```

Results

The administrative workstation migration is complete. Perform maintenance functions from the new administrative workstation only.

Chapter 2. Configuring databases

When your business needs change, you can create objects to extend the scope of your database before you create additional applications. As configuration changes are based on business rational, you must understand the structure of your database and the business requirements before you begin creating objects.

Database design

You typically extend the scope of your system because you have a use case that you want to include in the system. During the design phase you define the use case, the business objects, and the relationships between the business objects. To configure objects in the database, you must understand the structure of the database and the implications of the changes that you make to existing objects, tables, indexes, and relationships.

Relational database structure

The database and the database structure are defined in the installation process. The structure of the database depends on whether the database is Oracle Database, IBM DB2®, or Microsoft SQL Server.

A database that can be perceived as a set of tables and manipulated in accordance with the relational model of data. Each database includes:

- a set of system catalog tables that describe the logical and physical structure of the data
- a configuration file containing the parameter values allocated for the database
- a recovery log with ongoing transactions and archivable transactions

Table 3. Database hierarchy

Component	Description
<i>Data dictionary</i>	A repository of information about the application programs, databases, logical data models, and authorizations for an organization. When you change the data dictionary, the change process includes edit checks that can prevent the data dictionary from being corrupted. The only way to recover a data dictionary is to restore it from a backup.
<i>Container</i>	A data storage location, for example, a file, directory, or device that is used to define a database.
Storage partition	A logical unit of storage in a database such as a collection of containers. Database storage partitions are called <i>table spaces</i> in DB2 and Oracle, and called <i>file groups</i> in SQL Server.

Table 3. Database hierarchy (continued)

Component	Description
Business object	A tangible entity within an application that users create, access, and manipulates while performing a use case. Business objects within a system are typically stateful, persistent, and long-lived. Business objects contain business data and model the business behavior.
Database object	An object that exists in an installation of a database system, such as an instance, a database, a database partition group, a buffer pool, a table, or an index. A database object holds data and has no behavior.
Table	A database object that holds a collection of data for a specific topic. Tables consist of rows and columns.
Column	The vertical component of a database table. A column has a name and a particular data type for example, character, decimal, or integer.
Row	The horizontal component of a table, consisting of a sequence of values, one for each column of the table.
View	A logical table that is based on data stored in an underlying set of tables. The data returned by a view is determined by a SELECT statement that is run on the underlying tables.
Index	A set of pointers that is logically ordered by the values of a key. Indexes provide quick access to data and can enforce uniqueness of the key values for the rows in the table.
Relationship	A link between one or more objects that is created by specifying a join statement.
Join	An SQL relational operation in which data can be retrieved from two tables, typically based on a join condition specifying join columns.

Data dictionary tables

The structure of a relational database is stored in the data dictionary tables of the database.

The following table describes the product data dictionary tables.

Table 4. Tables in the product data dictionary

Table name	Contents
MAXOBJECT	All objects. Links an object to its table or view.
MAXTABLE	All tables.
MAXVIEW	All views.

Table 4. Tables in the product data dictionary (continued)

Table name	Contents
MAXATTRIBUTE	All attributes of an object. A table or view attribute depends on the attributes of the object.
MAXVIEWCOLUMN	All view columns.
MAXRELATIONSHIP	All relationships defined on objects.
MAXSEQUENCE	All sequences used in the system. In SQL Server, the sequences are generated from this table. Oracle and DB2 use database sequence generators.
MAXSYSINDEXES	All indexes in the system. This table contains the index name, uniqueness and storage partition that is used in the MAXSYSKEYS table.
MAXSYSKEYS	The columns in an index.

Integrity checker

The integrity checker is a database configuration utility that you can use to assesses the health of the base layer data dictionary. The tool compares the data dictionary with the underlying physical database schema. If errors are detected, the tool produces error messages detailing how to resolve the issues.

You run the integrity checker in the source environment before and after you upgrade the database. Activities that might result in errors include:

- Running the upgrade process itself
- Running the set of patch scripts included in updating a database
- Configuring the database in the Database Configuration application
- Configuring the database in the Migration Manager application

Errors reported by the integrity checker might affect migration. You run the integrity checker in both the source and the target environments when you run the Migration Manager.

You must ensure that the errors are corrected either by using the integrity checker in repair mode or by applying changes directly to the underlying database.

Storage partitions

A database storage partition is the location where a database object is stored on a disk. Database storage partitions are called *table spaces* in DB2 and Oracle, and called *file groups* in SQL Server.

When a database is created, the database administrator configures the DBSTORAGEPARTITION domain in the Domains application to include a list of available table spaces where objects are stored. When you create an object as a table, you specify the storage partition from this list of available table spaces.

In IBM DB2, the database or system can manage table spaces:

- If the database manages table spaces, indexes can be different from the table.
- If the system manages table spaces, indexes must be the same as the table.

The database and the system cannot manage table spaces simultaneously; you must choose one or the other.

Business objects

A business object is an object that has a set of attributes and values, operations, and relationships to other business objects. Business objects contain business data and model the business behaviour.

Unlike a business object, a database object does not model behavior. A database object is a self-contained software entity that consists of both data and functions to manipulate data. A business object might comprise of one or more database objects.

Every business object has a fixed set of properties that identify the business object type. The properties also specify how the database can use the business object.

Information about a business object, which is referred to as the metadata, is stored in the database in database tables. Business objects include the following metadata:

- The definition of the business object, such as its name, the database entity, whether the object is persistent or nonpersistent, and the Java class name
- Attributes, such as the name, data type, size, and the field validation class name
- Associated relationships

There are two types of business objects: persistent business objects and nonpersistent business objects. A persistent business object stores attribute values in a database. A nonpersistent business object does not store any metadata. Data in a nonpersistent business object is transient and is never stored in the database.

The metadata for a persistent business object represents the data in a database table or view.

The metadata that is associated with the business objects is used to manage the database objects. As a result, a database table or view is always required to be associated with a persistent business object.

Related tasks:

“Creating objects” on page 56

An object is a self-contained software entity that consists of both data and functions to manipulate data. You can use an abbreviation of your organization as a prefix to any new or modified object or attribute name.

“Adding attributes to objects” on page 59

You can add attributes to an object when you want to provide more information about that object.

“Changing attributes” on page 60

You can change attributes when necessary. Depending on how the attribute is configured, you might not be able to modify all of the fields.

“Creating restrictions on attributes” on page 61

You can create restrictions on attributes to prevent external data from overwriting the value for the chosen attribute.

“Defining lookup maps” on page 78

You define a lookup map to associate a source object and a source field with a target object and related fields. The product is delivered with pre-defined lookup maps, but you can define your own lookup maps for objects that you create.

Related reference:

“Attribute data types” on page 53

Each database record contains multiple attributes. Every attribute has an associated data type.

User-defined objects

Objects can be created in two ways: you can create an object in the database or an object can be natively defined in the database. User-defined objects are always created in the Database Configuration application.

Existing or imported objects are first natively defined in the backend of the database. They are later redefined in Maximo in the Database Configuration application. When an object is imported into the database, the **Imported** check box is automatically selected on the **Objects** tab.

Configuration levels for objects

Levels describe the scope of objects and must be applied to objects. Depending on the level that you assign to objects, you must create certain attributes. For users to access an object, an attribute value must exist at the level to which they have authority. The level that you assign to an object sometimes depends on the level of the record in the database.

A system-level object is the only object that does not require an attribute value as it applies to all objects. If you specify a multiple-level for an object, multiple attributes must be created. For example, if you specify the SYSTEMORGSITE level, the system attribute, the organization attribute, and the site attribute must be created.

You do not need to specify the required values for attributes when you create an object, required values can be specified at a later date.

Security is applied to configuration levels.

For certain configuration levels, you can restrict the result set by appending a condition to the WHERE clause. For example, you can specify the site level as "siteid=...".

Table 5. Configuration levels for database objects

Level	Description	Object attributes	Example
SYSTEM	A system-level object. Security restrictions are applied at the application or object level in the specific system-level business object definitions.	System attribute	
SYSTEMORG	A system-level object that can also be assigned to an organization. If the organization ID is not specified, the object operates at the system level.	System attribute and organization attribute	orgid is null or orgid = ...
SYSTEMSITE	A system-level object that can also be assigned to a site. If the site ID is not specified, the object operates at the system level.	System attribute and site attribute	siteid is null or siteid = ...
SYSTEMORGSITE	A system-level object that can also be assigned to an organization, or to an organization and a site. If the site ID is not specified, the object operates at either the system level or the organization level. The level depends on whether the organization ID is assigned. If the organization ID is not specified, the object operates at the system level.	System attribute, organization attribute, and site attribute	(siteid is null or siteid = ...) and (orgid is null or orgid = ...)

Table 5. Configuration levels for database objects (continued)

Level	Description	Object attributes	Example
SYSTEMAPPFILTER	This object is treated as a system-level object but it can ask the profile for a list of sites and organizations in the context of an application so that the application can filter data. Filtering is required for site-level administration of users and groups. Used for Users and Groups.	System attribute and application filter attribute	
ORG	An organization-level object. The framework applies security for this type.	Organization attribute	orgid = ...
ORGSITE	An organization-level object that can also be assigned to a site. If the site ID is not specified, the object operates at the organization level.	Organization attribute and site attribute	(siteid is null or siteid = ...) and orgid = ...
ORGAPPFILTER	An organization-level object with application filtering. Used for contracts so that the contract applications can filter on the special object instead of filtering by using standard security.	Organization attribute and application filter attribute	
SITE	A site level object.	Site attribute	siteid = ...
SITEAPPFILTER	A site-level object with application filtering. Reserved for future objects.	Site attribute and application filter attribute	

Table 5. Configuration levels for database objects (continued)

Level	Description	Object attributes	Example
ITEMSET	An item set-level object. The itemsetid attribute value must exist in the insert organization for users. The framework adds the required security restriction.	Item set attribute	
COMPANYSET	A company set-level object. The compnaysetid attribute value must exist in the insert organization for users. The framework adds the required security restriction.	Company set attribute	

Database relationships

Database relationships are associations between tables that are created using join statements to retrieve data.

The following table describes the database relationships.

Table 6. Database relationships

Type of relationship	Description
One-to-one	Both tables can have only one record on each side of the relationship. Each primary key value relates to none or only one record in the related table. Most one-to-one relationships are forced by business rules and do not flow naturally from the data. Without such a rule, you can typically combine both tables without breaking any normalization rules.
One-to-many	The primary key table contains only one record that relates to none, one, or many records in the related table.
Many-to-many	Each record in both tables can relate to none or any number of records in the other table. These relationships require a third table, called an associate or linking table, because relational systems cannot directly accommodate the relationship.

In the Database Configuration application, you can define Structured Query Language (SQL) statements for joins, and create relationships between parent and child objects. You can use a join to link data from multiple objects. The parent is the existing object and the child is the object that you are create.

Example

Parent = MAXUSER, Child = SITE, and Name = DEFSITE means that maxuser exists and you want to get the site for the default site for the user.

```
siteid = :defsite
```

This configuration means `site.siteid = maxuser.defsite`. When the SQL statement is run, the value of the parent attribute replaces anything preceded by a colon.

Business object attributes

Attributes of business objects contain the data that is associated with a business object. A persistent attribute represents a database table column or a database view column. A nonpersistent attribute exists in memory only, because the data that is associated with the attribute is not stored in the database.

A persistent business object can have persistent and nonpersistent attributes. Persistent attributes of a business object relate to columns of a database table or view. All attributes of a non-persistent business object are non-persistent.

The additional metadata that is associated with business object attributes is stored apart from the basic data type information. For example, attributes can include a domain, a custom class, a default value, and to specify whether the attribute is required.

Restrictions on attributes

Before you modify an attribute, you can verify whether it was created by the system or by someone at your site. Attributes created by the system have more restrictions on modifications than user-defined attributes. You cannot delete attributes created by the system.

In integration scenarios, data for a business object might be received from external business applications. Restricting changes to attributes prevents external data from overwriting the value of an attribute.

Some restrictions depend on whether text search is enabled for the object or on the data type. The rules governing modifications vary by attribute. For example, certain data types have a set value for the length, scale, dates, or integers. The **Memo** field is a regular ALN data type and it does not contain restricted values.

To manage restrictions on attributes, you can perform the following actions:

- View the current object attribute restrictions
- Restrict the attributes of an object
- Remove attribute restrictions from an object

Attribute data types

Each database record contains multiple attributes. Every attribute has an associated data type.

Table 7. Attribute data types

Data type	Data type name	Description
ALN	Alphanumeric characters, mixed case	Maximum length depends on the database: • Oracle = 4000 characters • SQL Server = 8000 characters • DB2 = 32672 characters
AMOUNT	Decimal number, used for currency	
BIGINT	Big integer	
BLOB	Binary large object	Stores JPEG, movies, or PDF files in single records inside the database instead of in external files.
CLOB	Character large object	
CRYPTO	Encrypted binary	Encrypts data on the screen and in the database. Used for password hints.
CRYPTOX	Encrypted binary (one-way)	Encrypts data in the database, but leaves it readable on the screen. Used for passwords.
DATE	Date only	
DATETIME	Date and time	
DECIMAL	Decimal number	A number that includes an integer and a fraction that consists of a fixed number of digits called the scale.
DURATION	Duration in hours	Appears as 1:30 = 1.5 hours
FLOAT	Floating number	Numbers with fractional portions with variable precision.
GL	General ledger account	An ALN data type that is used for GL Accounts.
INTEGER	Integer number	
LONGALN	Long alphanumeric.	Used only for nonpersistent Long Description attributes. The corresponding native column in the database is defined as CLOB.
LOWER	Lowercase characters	
SMALLINT	Small integer	
TIME	Time only	
UPPER	Uppercase characters	
VARCHAR	Variable length character	
YORN	Yes or No, 1 or 0 in the database	

Related concepts:

“Business objects” on page 48

A business object is an object that has a set of attributes and values, operations, and relationships to other business objects. Business objects contain business data and model the business behaviour.

“Storage partitions” on page 47

A database storage partition is the location where a database object is stored on a disk. Database storage partitions are called *table spaces* in DB2 and Oracle, and called *file groups* in SQL Server.

Related tasks:

“Creating objects” on page 56

An object is a self-contained software entity that consists of both data and functions to manipulate data. You can use an abbreviation of your organization as a prefix to any new or modified object or attribute name.

“Adding attributes to objects” on page 59

You can add attributes to an object when you want to provide more information about that object.

“Changing attributes” on page 60

You can change attributes when necessary. Depending on how the attribute is configured, you might not be able to modify all of the fields.

“Creating restrictions on attributes” on page 61

You can create restrictions on attributes to prevent external data from overwriting the value for the chosen attribute.

“Defining lookup maps” on page 78

You define a lookup map to associate a source object and a source field with a target object and related fields. The product is delivered with pre-defined lookup maps, but you can define your own lookup maps for objects that you create.

Views

A view does not contain data. The view is a definition that sits in the data dictionary with a database query that retrieves its data. Therefore, a view can contain data from more than one object, row, or attribute. When you fetch the data from a view, the database pulls the necessary records based on the WHERE clause and returns the data. Attributes are loaded when you create a view object.

A view is populated depending on the object on which it is based. For example, if a view is based on the WORKORDER object, the **Automatically Select** check box is selected, and you add or remove an attribute to this object, the attribute is either added or removed from the view. When you change an attribute, not all changes are applied to the associated view. For example, if you change the data type of an attribute, the change is applied to the view. However, if you change or add a domain to the default value of the WORKORDER object, the change is not applied to the view. Instead, you must apply this change to the view.

Because views are stored as named queries, you can use views to store frequently used, complex queries. You can run the queries using the name of the view in a simple query. **View Select** is optional. When **View Select** is not specified, all columns in both tables are included.

Tables can contain many columns and rows. To filter data, you can create a view that is a subset of a database that an application can process. The view can contain parts of one or more tables.

Sometimes you want to consider only parts of a table or parts of several tables. Data samples might include:

- Some of the columns.
- Rows that satisfy a certain condition.
- Some columns of one table and some columns of a related table.

Indexes

You can use indexes to optimize performance for fetching data. Indexes provide pointers to locations of frequently accessed data. You can create an index on the columns in an object that you frequently query.

You cannot redefine existing indexes. You must delete indexes and re-create their definitions.

Defining objects for applications

To extend the scope of your system, you must create objects. Every application is created based on the underlying objects that support the required business functions. You must understand the construction of your database before you define additional objects.

Creating objects

An object is a self-contained software entity that consists of both data and functions to manipulate data. You can use an abbreviation of your organization as a prefix to any new or modified object or attribute name.

Procedure

1. In the Database Configuration application, click **New Object**.
2. Type a name in the **Object** field. The **Entity** field shows the value that you typed in the **Object** field and becomes the name of the view on the database.
3. Specify a description for the object.
4. Optional: Select the details for the objects:

Option	Description
Main Object	To make the object a main object for Workflow, select this check box.
Persistent	If the object is persistent, three attributes are created: ID, description, and rowstamp (if selected). If the object is nonpersistent, nothing is added for attributes. You cannot configure the database without creating at least one attribute for the object.
Storage Partition	If applicable to your database, click Detail and select a storage partition for the object. The values are stored in the DBSTORAGEPARTITION domain.
User Defined	If an administrator created the object, the User Defined check box is selected. If the object is a regular product object, the User Defined check box is cleared.

Option	Description
Unique Column	The name of the attribute that is created as a unique identifier on a persistent object. This value is used in indexing. • If the object is flagged as imported, then a unique column is not required. • If you add a unique column, it must have a new column name and cannot exist in the native database.
Language Table	To enable this object for multiple languages, specify a value. The convention is <i>L_tablename</i> .
Text Search Enabled	Select to enable text search on the object. You can use this function with text search on attributes.

5. Click **Save Object**.

What to do next

After you create an object, you must add attributes before you configure the database for your changes to take effect.

Related concepts:

“Business objects” on page 48

A business object is an object that has a set of attributes and values, operations, and relationships to other business objects. Business objects contain business data and model the business behaviour.

“Storage partitions” on page 47

A database storage partition is the location where a database object is stored on a disk. Database storage partitions are called *table spaces* in DB2 and Oracle, and called *file groups* in SQL Server.

Related reference:

“Attribute data types” on page 53

Each database record contains multiple attributes. Every attribute has an associated data type.

Adding views to databases

A view can contain data from more than one object in the database. You can create object views in addition to views on existing system objects. If you work with custom applications, you can add object tables in the Database Configuration application.

About this task

You can add an object table or view from any tab of the Database Configuration application.

Procedure

1. On the toolbar, click **New Object**, and specify a value and a description for the object.
2. Optional: Type over the value in the **Entity** field to change the name of the object on the native database.

3. In the **Service** field, specify a value or use the default value of CUSTAPP.
4. In the **Level** field, specify the scope of the object in the Multisite scheme.
5. Optional: To create a view, select an object in the **Extends Object** field. The **View** check box is selected by default.
6. Optional: Provide the following additional object details:

Option	Description
Main Object	To make the object a main object for Workflow, select this check box.
Persistent	If the object is persistent, the check box is selected and three attributes are created: ID, description, and rowstamp. If the object is non-persistent, the check box is clear. Nothing is added for attributes but you cannot configure the database without creating at least one attribute for the object.
User Defined	If the object is a regular product object, the User Defined check box is clear. If the object was created by an administrator the User Defined check box is selected.
Storage Partition	If applicable to your database, specify a storage partition for the object.
Unique Column	The name of the attribute that is created as a unique identifier on a persistent object. This value is used in indexing. • If the object is flagged as imported, then a unique column is not required. • If you add a unique column, it must have a new column name and cannot exist in the native database.
Text Search Enabled	To enable text search on the object, select the check box. You can use this function with text search on attributes. (This field appears only for existing objects.)

7. Optional: In the **View** section, define the following details for a view:

Option	Description
View Where	The WHERE clause that is used for the view.
Join to Object	The secondary object that is used in the join for this view. If the view joins two tables, you can type the name for the second table in this field.
View Select	The SELECT clause that is used for the view when the Automatically Select check box is cleared. Use the format SELECT TABLE1.COL1 AS A, TABLE2.COL2 AS B
View From	The FROM clause that is used for the view when the Automatically Select check box is cleared.

8. Optional: In the Audit table window, create an audit table and select the **Audit Enabled** check box to edit the filter field for the E-audit function.
9. Save the object.

What to do next

After you add an object table or view, you must configure the database for your changes to take effect.

Specifying attributes for objects

Attributes contain the data that is associated with objects. When you create an object, you must specify the attributes that are assigned to the object.

Descriptions and long descriptions

There is a limit to the amount of data that you can store in the description fields. Each database has its own specified limit. If you know that you need to store more data than is available in the description field, you can add a long description field to an attribute.

The description field object is associated with the VARCHAR (variable length character) column. This column can hold a maximum number of bytes it can hold. If you associate a long description field with an attribute, a long description is added to the VARCHAR column. Long descriptions are stored in the character large object (CLOB) column of the long description (LONGDESCRIPTION) table. If you include a long description field, you can store a smaller amount of data in the main table and a larger amount in the long description table.

The long description data can contain any alphanumeric data. This data appears in the Long Description field from the user interface.

Each long description must have a long description owner. When you select the **Long Description Owner** check box on the **Attributes** tab of the Database Configuration application, you become the owner of the long description. When the **Long Description Owner** check box is selected for an attribute for the first time, the HASLD (has long description) column is added to the main record.

An icon exists or can be added beside the description field on the UI. The appearance of the icon changes when long description data exists. HASLD indicates that a long description exists.

The default value for HASLD is 0. When data is entered in the long description, the value of HASLD changes to 1.

Adding attributes to objects

You can add attributes to an object when you want to provide more information about that object.

About this task

You can use an abbreviation of your organization as a prefix to any attribute name, for example, ACME_MEMOFIELD. This practice prevents accidentally choosing a database reserved word and prevents conflicts with new standard names in an upgrade.

New attributes are accessible in the user interface by first adding them to the respective application by using the Application Designer function. You can then add the attributes to the respective application in the Application Designer application to make the attributes accessible on the user interface.

Procedure

1. On the **List** tab, select the object to which you want to add an attribute.
2. On the **Attributes** tab, click **New Row**.
3. In the **Attribute** field, specify a value.
4. Specify values in the **Title** field and the **Description** field.
5. Optional: If you want to include a long description which can contain more data than the **Description** field, select the **Long Description Owner** check box.
6. In the **Type** field, specify the data type of the attribute. The **Length** and **Scale** fields show default values based on the type that you select.
7. Optional: Select the **Required** check box if you want this attribute to represent a required field on the screen. If this is a persistent attribute it is also required on the database.

Requirement: If data exists in the table to which the attribute belongs, and you select the **Required** check box, you must specify a value in the **Default Value** field. Specifying a value in this field ensures that you cannot enter a null value for the required attribute. The default value you specify is only validated against the associated domain for the attribute and the data type of the attribute.

8. Optional: If you want a group of attributes to share the same data type and length, specify the parent attribute in the **Same as attribute** field. The child attribute copies the details of the parent attribute. If you change the details of any attribute in the group, the details of all attributes in the group change.
9. Optional: Provide additional information about the attribute and save the object.

What to do next

You must configure the database for your changes to take effect.

Related concepts:

“Business objects” on page 48

A business object is an object that has a set of attributes and values, operations, and relationships to other business objects. Business objects contain business data and model the business behaviour.

“Storage partitions” on page 47

A database storage partition is the location where a database object is stored on a disk. Database storage partitions are called *table spaces* in DB2 and Oracle, and called *file groups* in SQL Server.

Related reference:

“Attribute data types” on page 53

Each database record contains multiple attributes. Every attribute has an associated data type.

Changing attributes

You can change attributes when necessary. Depending on how the attribute is configured, you might not be able to modify all of the fields.

Before you begin

Verify whether the attribute was created by the system or by someone at your site. If it was created by someone at your site, the **User Defined** check box is selected. You cannot delete attributes created by the system.

About this task

Attributes created by the system have more restrictions on modifications than user-defined attributes do. Some restrictions depend on whether text search is enabled for the object or on the data type.

For example, certain data types have a set value for the length, scale, dates, or integers. The **Memo** field is a regular ALN. You can make it anything you want.

The rules governing modifications are complex, and vary by attribute.

Procedure

1. Locate the attribute that you want to modify.

Tip: You can use **Advanced Search** to search for attributes.

2. Edit the **Description**, **Type**, **Length**, and **Required** fields according to your business needs. Some fields are read only, depending on values in other areas.
3. Click **Save Object**. The status of all affected objects display To Be Changed until you configure the database.

What to do next

You must configure the database for your changes to take effect.

Related concepts:

“Business objects” on page 48

A business object is an object that has a set of attributes and values, operations, and relationships to other business objects. Business objects contain business data and model the business behaviour.

“Storage partitions” on page 47

A database storage partition is the location where a database object is stored on a disk. Database storage partitions are called *table spaces* in DB2 and Oracle, and called *file groups* in SQL Server.

Related reference:

“Attribute data types” on page 53

Each database record contains multiple attributes. Every attribute has an associated data type.

Creating restrictions on attributes

You can create restrictions on attributes to prevent external data from overwriting the value for the chosen attribute.

Procedure

1. In the Database Configuration application, use the **List** tab to search for and to locate the object whose attributes you want to restrict.
2. Click the **Object** tab to make the selected object the current object in the application.
3. From the **Select Action** menu, select **Restrict Attributes**.

4. Navigate to the attribute that you want to restrict.

Tip: You can use the **Filter** fields to help you locate the attribute.

5. Select the **Restricted** check box.
6. Click **OK**.

What to do next

You must configure the database for your changes to take effect.

Related concepts:

“Business objects” on page 48

A business object is an object that has a set of attributes and values, operations, and relationships to other business objects. Business objects contain business data and model the business behaviour.

“Storage partitions” on page 47

A database storage partition is the location where a database object is stored on a disk. Database storage partitions are called *table spaces* in DB2 and Oracle, and called *file groups* in SQL Server.

Related reference:

“Attribute data types” on page 53

Each database record contains multiple attributes. Every attribute has an associated data type.

Excluding user-defined attributes when duplicating objects

There might be objects that you duplicate regularly but you do not want to include all user-defined attributes in the duplicated objects. For each object, you can specify the user-defined attributes that you do not want included when you duplicate objects. By excluding attributes, you can reduce the data maintenance work required to work with duplicated objects.

About this task

The list of attributes available for selection in the Skip Attributes window is populated by the user-defined attributes that are available for the selected object. If user-defined attributes do not exist in the object, no attributes are available for you to select. When you select user-defined attributes to be skipped, the attributes are not listed in future selection options.

If you specify that user-defined attributes be skipped for base objects, the skip values also apply to the object views.

Procedure

1. In the Database Configuration application, select an object that you want to duplicate.
2. From the **Select Action** menu, select **Skip Attributes**.
3. Specify the attributes that you do not want to be copied to new objects.
4. Optional: To specify conditions when attributes are skipped, specify a value in the **Condition** field.
5. Click **OK** and save the object.

What to do next

If you specify attributes to be skipped in the duplication process, you must configure the database for the changes to take effect.

Enabling autonumbering for attributes

You can configure autonumbering at the system, set, organization, or site level. You can enable autonumbering for attributes to apply starting numbers and prefixes to items such as assets or work orders. Autonumbering can facilitate the movement of items across the organization levels.

Before you begin

You can specify an autonumber only if the **Can Autonumber** check box is selected on the **Attributes** tab of the Database Configuration application. Whether this check box is selected depends on the data type of the attribute.

Procedure

1. In the Database Configuration application, select an object.
2. On the **Attributes** tab, select an attribute.
3. In the Advanced section, specify an autonumber. You can create an autonumber or use an existing autonumber.
4. In the Details section, specify &AUTOKEY& as the default value for the attribute and save the object.
5. If you created an autonumber in step 3, use the **Autonumber Setup** action in the Organizations application to specify the seed value and the prefix.
6. Save the object.

What to do next

You must configure the database for your changes to take effect.

Adding tax types to database tables

A tax type corresponds to a kind of tax, for example, to a city sales tax. You can specify up to 27 different tax types. The requirements of your financial system determine how many tax types you can specify.

Before you begin

Administration mode must be turned on before you can configure the database.

About this task

When you add tax types to the database, you specify the number of tax types, you configure the database, and then you update the database tables for your organizations. If you add more than five tax types, and configure and update the database tables, all of the tax types display in the Tax Options window of the Organizations application.

Procedure

1. In the Database Configuration application, from the **Select Action** menu, select **Tax Types > Add/Modify Tax Types**.

2. In the **Number of Tax Types** field, specify the number of tax types that you want to use.
3. Click **OK**.
4. From the **Select Action** menu, select **Apply Configuration Changes** to configure the database.
5. In the Database Configuration window, provide the required information, then click **OK**.
6. From the **Select Action** menu, select **Tax Types > Update Tax Data**, to update the database tables with the changes.

Adding indexes

You create an index to optimize performance when searching the table using the **Find** function and to establish uniqueness of table columns.

About this task

Indexes can be defined only for persistent tables.

Procedure

1. From the **List** tab, select the object to which you want to add an index, and click the **Indexes** tab.
2. In the Indexes table, click **New Row**.
3. In the **Index** field, specify a value.
4. Optional: To make each column in the index unique, select the **Enforce Uniqueness** check box.
5. For IBM DB2 and SQL Server: Select the **Clustered Index** check box to create a clustered index. You can have only one clustered index per table.
6. Specify a storage partition or accept the default value.
7. In the Columns table, click **New Row** to add a column to the index.
 - a. In the **Column** field, specify an attribute from the selected object.
 - b. Optional: In the **Sequence** field, specify the sequence for the column. If you do not specify a value, then the order in which you add columns determines their sequence.
 - c. Optional: Select the **Ascending** check box to have the index searched in ascending order. If this field is cleared, the index is searched in descending order.
8. Click **Save Object**. The index status will display Add until you configure the database.

What to do next

You must configure the database for your changes to take effect.

Related concepts:

“Indexes” on page 56

You can use indexes to optimize performance for fetching data. Indexes provide pointers to locations of frequently accessed data. You can create an index on the columns in an object that you frequently query.

Creating relationships between parent and child tables

To retrieve data on objects, you must define the relationships between objects. You define SQL for joins to create relationships between parent and child objects.

Before you begin

Before you create a relationship, review the relationships defined in the database to determine if you can reuse an existing relationship. If the definition of a relationship does not match your needs, do not modify the relationship. Instead, create an additional relationship.

About this task

You define SQL for *JOINS*. A JOIN lets you link together data from multiple tables.

Procedure

1. Use the **List** tab to select the table for which you want to create a relationship, and then click the **Relationships** tab.
2. Click **New Row**.
3. In the **Relationship** field, specify a name.
4. Create a WHERE clause.
5. Select a **Child Object**.
6. Optional: Type comments in the **Remarks** field.
7. Click **Save Object**.

What to do next

If the relationship that you created is for a table or attribute, the changes will not take effect until you configure the database.

Related concepts:

“Database relationships” on page 52

Database relationships are associations between tables that are created using join statements to retrieve data.

Account format configuration for General Ledger

Default General Ledger (GL) accounts are created as part of the installation process but you can create additional GL account codes at a later date. You must configure default GL account code formats for the system level before you can create Chart of Accounts. Once configured, it is not recommended that you modify the GL account formats as changing the configuration can make the existing Chart of Accounts unusable.

General Ledger account configuration

Each General Ledger account code consists of several components (segments). In the Database Configuration application, you define the default, system-level format of the account code. In the Chart of Accounts application, you specify the valid components to be used, as well as the organization-specific format of the account code.

For easy identification, use delimiters to separate components when they display. For example, use hyphens to separate components: 6100-400-SAF. Delimiters are written to the database.

For any account code, you can:

- Define 20 or fewer components.
- Restrict the number of characters in a component field.

- Include a total of 254 or fewer characters/digits.

You can use any tab of the Database Configuration application to specify the general ledger account code format.

Related reference:

“General Ledger account components”

Some General Ledger account components are required while other components are optional.

General Ledger account components

Some General Ledger account components are required while other components are optional.

Table 8. Required and optional components

Type of component	On-screen display
Required	Unknown values not specific to required components contain placeholder characters.
Optional	Any unknown optional components do not display. In the DEMO database, the fourth component is optional (most account codes consist of the first three components). • It does not require any characters. • No accounts have been assigned to it in Chart of Accounts, so it does not appear as part of the General Ledger Account.

Your General Ledger system has rules regarding whether an account is acceptable when partially defined.

- Fully defined (fully specified) accounts
 - Have no unknown values (placeholders) in required components
 - Example: 6100-350-SAF is fully defined
- Partially defined (partially specified) accounts
 - Contain placeholders in some required components
 - Example: 6100-???-SAF (the required Activity component is not specified and therefore contains placeholder characters)

General Ledger component sequence

Account components display in a sequential format, with the first component in the string representing the highest level.

Example

For example, the MAXDEMO database includes the following component sequence:

- Component 1 is for cost center
- Component 2 is for activity
- Component 3 is for resource
- Component 4 is for element

Since account components are concatenated, with the highest level at the beginning, account 6100-350-SAF is represented as:

Table 9. General Ledger account component sequence

Component 1	Component 2	Component 3	Component 4
6100	350	SAF	(not used)
Cost center	Activity	Resource	Element

General Ledger component lengths

Changing the length of the component values can result in invalid general ledger accounts. If you change the length, change the values to fit the new length.

Example

In maxdemo the cost center component length is 4, the resource and activity component lengths are both 3, and the element component is 10. When you add in the three delimiters, the length of the GL is 23.

If you change the cost center component length to 3 and the activity component length to 4, the total length remains 23. No configuration is required. However, the general ledger component is now invalid. The cost center component length was shortened to 3 but has a four-digit value (in this example) of 6000.

Configuring the database

When you change the database, for example by creating or deleting objects, attributes, or indexes, changes are stored in secondary tables. The changes do not take effect until you configure the database. The product restores the backup tables as part of configuration.

Modes of configuring the database

There are three ways to configure the database: command-line mode, a partial live configuration, or a full live configuration with administration mode turned on. The option that you chose depends on whether the application server needs to be shut down and the impact on access to users.

The following table lists the advantages and disadvantages of each configuration mode:

Table 10. Configuration modes

Configuration modes	Advantages	Disadvantages
Command-line mode	All data, logs, and control files are backed-up. It is easier to restore data from this configuration mode as there are no updates while the database is being configured.	Must shut down the application server therefore, users have no access to the applications during the configuration. Requires an information technology administration user who has control over the application server.

Table 10. Configuration modes (continued)

Configuration modes	Advantages	Disadvantages
Full live with administration mode turned on	Allows an administration user to perform tasks such as adding a column to a table without shutting down the application server. Users have access to applications unrelated to the configuration changes.	Blocks users from the system applications. Suspends CRON tasks. Does not allow remote connectivity Disables event listeners. Requires the user to have Administrator login security authorizations, which you must assign in the Security Groups application.
Partial live	Least impact on users as active transactions are not interrupted or lost. Does not require administration mode to be turned on. If changes performed during a live update are later determined to be incorrect, you can perform another live update to reverse the changes. Changes do not disrupt the live business object definitions. If you change the field validation class for an attribute and perform a live update to apply the changes, the business objects that are already instantiated are not revalidated.	Can be used only for changes with Maximo that have no impact on users such as changing an attribute name.

Related tasks:

“Configuring the database in command-line mode”

You can use command-line mode to configure the database.

“Configuring the database in administration mode” on page 69

You can configure the database in admin mode.

“Restoring backup tables” on page 70

You might need to restore your backup tables separately if you did not restore your backup tables during a command-line configuration.

Configuring the database in command-line mode

You can use command-line mode to configure the database.

About this task

When you save a record the changes are stored in temporary database configuration objects. Before you configure the database, you can close and reopen the Database Configuration application without losing any saved changes. A secondary table stores pending changes, which also appear in the **Status** field. You cannot query on **Status**.

Procedure

1. Shut down your application server and wait for one minute. The application server session timestamp updates every 60 seconds.
2. Open a command prompt and change the directory to *install_home\tools\maximo*.
3. To configure the database and restore the backup tables, type **configdb**. If you do not need to restore the backup tables, edit the *configdb.bat* file.
Sometimes the data in the temp tables (XX+tablename) must be modified before restoring. Perform the following steps:
 - a. Remove the -r parameter from the *configdb.bat* file.
 - b. Save your changes.
 - c. Return to the command prompt and type **configdb**.
4. If configuration errors occur, resolve the errors in the database. Open the log files for troubleshooting at *install_home\tools\maximo\log*.
5. Restart the application server. If backup tables were created, delete them before reconfiguring the database. If you ran **configdb** without the -r parameter, and if tables were rebuilt, you must restore the backup tables.

Related concepts:

"Modes of configuring the database" on page 67

There are three ways to configure the database: command-line mode, a partial live configuration, or a full live configuration with administration mode turned on. The option that you chose depends on whether the application server needs to be shut down and the impact on access to users.

Configuring the database in administration mode

You can configure the database in admin mode.

Before you begin

You must have administration login security authorization. You obtain this authorization from the Security Groups application.

Before you can configure and turn on administration mode, you must set the **mail.smtp.host property** to ensure that you can receive scheduled reports.

About this task

When you save a record the changes are stored in temporary database configuration objects. Before you configure the database, you can close and reopen the Database Configuration application without losing any saved changes. A secondary table stores pending changes, which also appear in the **Status** field. You cannot query on **Status**.

Procedure

1. In the Database Configuration application, select **Manage Admin Mode** from the action menu
2. In the Turn Admin Mode ON window, modify the values in the **Number of Administrative Sessions Allowed** field and the **Number of Minutes for User Logout** field. The default value of each field is 5. If you modify these fields, click **Update Properties** for the parameters to take effect.
3. Click **Turn Admin Mode ON**.
4. In the Electronic Signature Authentication window, enter the appropriate value in the **Reason for Change** field.
5. Click **OK**. A window opens that indicates that the Admin Mode is starting.
6. Click **OK**.
7. Throughout the configuration process, click **Refresh Status** to view the messages that the configuration process writes in the Status window. If you decide to cancel the configuration, click **Cancel Admin Mode**.
8. From the **Select Action** menu, click **Apply Configuration Changes** to configure the database and restore backup tables. Wait until administration mode is turned on before performing this step.
9. To turn off Admin Mode, from the **Select Action** menu, click the **Admin Mode** action, and then click **Turn Admin Mode OFF**.

Related concepts:

“Modes of configuring the database” on page 67

There are three ways to configure the database: command-line mode, a partial live configuration, or a full live configuration with administration mode turned on. The option that you chose depends on whether the application server needs to be shut down and the impact on access to users.

Restoring backup tables

You might need to restore your backup tables separately if you did not restore your backup tables during a command-line configuration.

Procedure

1. Open a command prompt and change the directory to: `install_home\tools\maximo`
2. Run `restorefrombackup`.
3. Start the application server.

Related concepts:

“Modes of configuring the database” on page 67

There are three ways to configure the database: command-line mode, a partial live configuration, or a full live configuration with administration mode turned on. The option that you chose depends on whether the application server needs to be shut down and the impact on access to users.

Configuring the system for regulatory compliance

You can use the electronic signature authentication and electronic audit to specify functions in applications that are tracked for auditing purposes. Users must provide their credentials to change records and also must provide a reason for changing the records.

Electronic signatures and audit records

Electronic signatures and audit records provide an additional level of security control and auditing capability. You can enable electronic signature and electronic audit records independently of one another, however they are typically used together.

The electronic auditing function writes audit records to the database tables. The audit tables are configured for each business object that is enabled for auditing.

Using electronic signatures and audit records involves the following functions:

- Login tracking
- Electronic signature
- Electronic audit

The following table lists how the controls are maintained.

Table 11. Electronic signature and audit record control points

Method	Description
Login Tracking Found in the Security Groups and Users application under Security Controls.	Controls the number of allowed login attempts and displays the current login status of the user.
Electronic signatures	Requires that the person saving a record, changing a record, or accessing a specific action is the person who logged in.
Electronic audit records	Records and audits changes to records, keeps copies of the changes, and produces an audit trail.

You define electronic signatures and audit records at the system level. When you enable electronic signatures and audit records, they apply to all organizations and sites.

Related tasks:

“Creating a drop-down list for the Reason for Change field” on page 74

Electronic signatures are enforced by requiring users to complete fields in an Electronic Signature Authentication window. The Electronic Signature Authentication window includes a **Reason For Change** field.

“Adding values to the Reason For Change domain” on page 75

If you elect to restrict users to giving only a select set of reasons in the **Reason for Change** field, you need to add values to the Reason for Change domain.

Related reference:

“Electronic signature authentication” on page 73

When users perform actions for which electronic signature is enabled, they must authenticate by entering data in the required fields. Authentication must be successful before users can continue.

Login tracking

With login tracking, you specify the number of permitted user login attempts. When the number is exceeded, any further login attempts are blocked. Login tracking also tracks the number of login attempts for a user and the login status of a user.

You can use login tracking independently of electronic signature, but you must enable login tracking to use electronic signature.

Electronic signature feature

Electronic signature records the user name and full user name of users who change database records or who perform actions in an application. The modification to the record such as change, insert, update, or delete, known as the identifier, is also recorded.

The full user name corresponds to the **Displayname** attribute in the Person object. When you add a user, you must associate a Person record for example, two workers are named John Smith. Their full names are John Allen Smith and John B. Smith.

After you enable electronic signature for a database attribute, when you try to save a change in a field that uses this database attribute or an implicit save is performed, you must authenticate your user credentials. All authentication attempts are saved in the LOGINTRACKING object. Authentication must be successful before the system saves the application data.

After you enable electronic signature for an action, when you access the action you must authenticate before you leave that page and window. Authentication must be successful before you can continue with the selected action.

During authentication, the LOGINTRACKING object records the following items:

- User name (login ID).
- Full user name (the person's display name).
- Date and time of the attempt.
- Whether the authentication was successful.
- Application name where the electronic signature was invoked.
- Reason for the change (as typed on the Electronic Signature Authentication window).
- Unique transaction identifier.
- Key values columns for the record.

Electronic audit records

Each time users add, delete, or modify the value of an attribute using a system application and save the change, an audit record is written to the audit object corresponding to the regular database object. Electronic audit records must first be enabled for the database attribute.

The audit record includes the following items:

- A copy of the changed data
- The user name of the user who changed the data
- The identifier indicating whether the change was an insert, update, or delete
- The current date and time of the transaction
- The rowstamp
- The unique e-Audit transaction ID
- The unique e-Sig transaction ID if electronic signature is enabled

- The key values columns for the record, even if those columns are not e- Audit enabled. For example, the work order number is recorded even when another attribute in the WORKORDER object triggers the electronic audit.

Electronic signature authentication

When users perform actions for which electronic signature is enabled, they must authenticate by entering data in the required fields. Authentication must be successful before users can continue.

The Electronic Signature Authentication window includes the following fields:

Table 12. Fields in the Electronic Signature Authentication window

Field	Description	Importance
User Name	Login ID	Required
Full user name (unlabeled)	Data is taken from the DISPLAYNAME attribute in the PERSON object	Read-only
Password	Password	Required
Reason for Change	Enter less than or equal to 50 characters	Required

Related concepts:

“Electronic signatures and audit records” on page 71

Electronic signatures and audit records provide an additional level of security control and auditing capability. You can enable electronic signature and electronic audit records independently of one another, however they are typically used together.

“Login tracking” on page 71

With login tracking, you specify the number of permitted user login attempts. When the number is exceeded, any further login attempts are blocked. Login tracking also tracks the number of login attempts for a user and the login status of a user.

“Electronic signature feature” on page 72

Electronic signature records the user name and full user name of users who change database records or who perform actions in an application. The modification to the record such as change, insert, update, or delete, known as the identifier, is also recorded.

“Electronic audit records” on page 72

Each time users add, delete, or modify the value of an attribute using a system application and save the change, an audit record is written to the audit object corresponding to the regular database object. Electronic audit records must first be enabled for the database attribute.

Related tasks:

“Creating a drop-down list for the Reason for Change field”

Electronic signatures are enforced by requiring users to complete fields in an Electronic Signature Authentication window. The Electronic Signature Authentication window includes a **Reason For Change** field.

“Adding values to the Reason For Change domain” on page 75

If you elect to restrict users to giving only a select set of reasons in the **Reason for Change** field, you need to add values to the Reason for Change domain.

Creating a drop-down list for the Reason for Change field

Electronic signatures are enforced by requiring users to complete fields in an Electronic Signature Authentication window. The Electronic Signature Authentication window includes a **Reason For Change** field.

Procedure

1. To make the **Reason For Change** field let users type free-form text, no steps are required.
2. To make the **Reason For Change** field require users to choose from a user-defined value list, you must add values to the CHANGEREASON domain.

Related concepts:

“Electronic signatures and audit records” on page 71

Electronic signatures and audit records provide an additional level of security control and auditing capability. You can enable electronic signature and electronic audit records independently of one another, however they are typically used together.

“Login tracking” on page 71

With login tracking, you specify the number of permitted user login attempts. When the number is exceeded, any further login attempts are blocked. Login tracking also tracks the number of login attempts for a user and the login status of a user.

“Electronic signature feature” on page 72

Electronic signature records the user name and full user name of users who change database records or who perform actions in an application. The modification to the record such as change, insert, update, or delete, known as the identifier, is also recorded.

“Electronic audit records” on page 72

Each time users add, delete, or modify the value of an attribute using a system application and save the change, an audit record is written to the audit object corresponding to the regular database object. Electronic audit records must first be enabled for the database attribute.

Related reference:

“Electronic signature authentication” on page 73

When users perform actions for which electronic signature is enabled, they must authenticate by entering data in the required fields. Authentication must be successful before users can continue.

Adding values to the Reason For Change domain

If you elect to restrict users to giving only a select set of reasons in the **Reason for Change** field, you need to add values to the Reason for Change domain.

Before you begin

You do not assign this value list to a database object and attribute. The connection to the database for this value list is already present.

Procedure

1. Open the **Domains** application.
2. Open the CHANGEREASON domain.
3. Click **Edit Detail**.
4. In the ALN Domain window, click **New Row**.
5. In the **Value** field and **Description** field, enter a value that you want the user to see in the CHANGEREASON value list. For this value list only:
 - These values are not written to the database.
 - The value in the **Description** field is what users see when they use the list.For example, you want users to see a value list containing only Change to Record and Delete Record. In the **Value** field and **Description** field, type the following information:

Value field	Description field
CHANGE	Change to record

Value field	Description field
DELETE	Deleted record

6. Click **OK**.

Related concepts:

“Electronic signatures and audit records” on page 71

Electronic signatures and audit records provide an additional level of security control and auditing capability. You can enable electronic signature and electronic audit records independently of one another, however they are typically used together.

“Login tracking” on page 71

With login tracking, you specify the number of permitted user login attempts. When the number is exceeded, any further login attempts are blocked. Login tracking also tracks the number of login attempts for a user and the login status of a user.

“Electronic signature feature” on page 72

Electronic signature records the user name and full user name of users who change database records or who perform actions in an application. The modification to the record such as change, insert, update, or delete, known as the identifier, is also recorded.

“Electronic audit records” on page 72

Each time users add, delete, or modify the value of an attribute using a system application and save the change, an audit record is written to the audit object corresponding to the regular database object. Electronic audit records must first be enabled for the database attribute.

Related reference:

“Electronic signature authentication” on page 73

When users perform actions for which electronic signature is enabled, they must authenticate by entering data in the required fields. Authentication must be successful before users can continue.

Database changes unrelated to eAudit

Some object-level and attribute-level parameters are eligible for live update.

Changes to the following object-level parameters are eligible for live update:

Table 13. Object-level parameters that are eligible for live updates

Description	Source
Description	MaxObjectCfg.Description
Esig Filter	MaxObjectCfg.EsigFilter

Changes to certain attribute-level parameters are eligible for live update. The following table lists these attribute-level parameters.

Table 14. Attribute-level parameters that are eligible for live updates

Header	Source	Additional rules to qualify for live update
Description	MaxAttributeCfg.Remarks	None
Title	MaxAttributeCfg.Title	None
Domain	MaxAttributeCfg.DomainID	None

Table 14. Attribute-level parameters that are eligible for live updates (continued)

Header	Source	Additional rules to qualify for live update
Default Value	MaxAttributeCfg.DefaultValue	None
Search Type	MaxAttributeCfg.SearchType	For live update, you cannot change to or from Text Search (domainid = SEARCHTYPE, maxvalue = TEXT)
Esig Enabled	MaxAttributeCfg.EsigEnabled	None
Can Autonum	MaxAttributeCfg.CanAutonum	None
Autokey Name	MaxAttributeCfg.AutokeyName	None
Is Positive	MaxAttributeCfg.IsPositive	None
Field Validation Class	MaxAttributeCfg.Classname	For live update, the specified class must be accessible to the class loader

If any of the Cfg tables for an object or any of its attributes have changes that are not listed in the preceding tables, that object and its attributes are not eligible for live update.

Database changes involving eAudit

An audit table is an object that MaxTableCfg.EauditTbname references.

Example

If the Person table is audited and the audit table is named A_Person, two rows exist in MaxTable Cfg, as shown in the following example.

Table 15. eAudit Tables

Table name	Eaudit Tbname
PERSON	A_PERSON
A_PERSON	(null)

For the base table, changes to the audit-related, object-level parameters in the following table are eligible for live update. For an audit table, if the object you add (MaxObjectGfg.Changed = 1) and its base table are eligible for live update, then the audit table is also eligible.

The following table lists the object-level parameters that are eligible for live update.

Table 16. Object-level parameters that are eligible for live update

Description	Source	Additional rules to qualify for live update
Eaudit Filter	MaxObjectCfg.EauditFilter	None

Table 16. Object-level parameters that are eligible for live update (continued)

Description	Source	Additional rules to qualify for live update
Eaudit Enabled	MaxObjectCfg.EauditEnabled	No additional rules. Field validation in the Database Configuration application ensures that when EauditEnabled is turned on, there is a nonnull value for EauditTbname.
Eaudit Table Name	MaxTableCfg.EauditTbname	If either of the following conditions is true, live update for change of EauditTbname is supported: - EauditTbname is non-null and the object referenced by EauditTbname is a new object (MaxObjectCfg.Changed = I. - EauditTbname is being set to null (EauditEnabled is being turned on).

The following table lists changes to the audit-related, attribute-level parameters on the base table that are eligible for live update.

Table 17. Attribute-level parameters that are eligible for live update

Description	Source	Additional rules to qualify for live update
Description	MaxAttributeCfg.EauditEnabled	No additional rules. Turning auditing on or off for an attribute that is eligible for live update, and does not involve native changes to the audit table.

Controlling changes to objects

Depending on your users needs, you can add functions that will help users, for example creating lookups or creating system messages.

Defining lookup maps

You define a lookup map to associate a source object and a source field with a target object and related fields. The product is delivered with pre-defined lookup maps, but you can define your own lookup maps for objects that you create.

Procedure

1. On the **Attributes** tab, select the attribute for which you want to specify details.
2. Click **Edit Lookup Map**.
3. Click **New Row**.
4. Specify the following details for the target attribute:
 - In the **Target Attribute** field, select the attribute to which you want to return values.

- In the **Source Object** field, specify the source from which values are returned.
 - In the **Source Key** field, specify the key of the source object from which you want to return values.
 - In the **Sequence** field, specify a numeric value to determine the attributes that are to be set before other attributes. A lower numeric value has priority over a higher numeric value.
5. Optional: Select the **Allow Null** check box if you want the target attribute to accept null values.
 6. Click **OK**.

Related concepts:

“Business objects” on page 48

A business object is an object that has a set of attributes and values, operations, and relationships to other business objects. Business objects contain business data and model the business behaviour.

“Storage partitions” on page 47

A database storage partition is the location where a database object is stored on a disk. Database storage partitions are called *table spaces* in DB2 and Oracle, and called *file groups* in SQL Server.

Related reference:

“Attribute data types” on page 53

Each database record contains multiple attributes. Every attribute has an associated data type.

Adding system messages

You can add, change, and delete system messages that display when required to assist users as they use the application.

Procedure

1. In the Database Configuration application, from the **Select Action** menu, select **Messages**.
2. Click **New Row** and specify the message group to which you want the new message to belong.
3. In the **Message Key** field, type a unique string that you want to be associated with the message.
4. Specify the display method for the message.
5. In the **Message ID** field, specify the prefix and product to be associated with the message. Customer-generated messages typically use the BMXZZ prefix. The complete identifier is generated when you click the prefix.
6. Select the **Display ID** check box to display the message identifier whenever the message is displayed.
7. In the **Value** field, enter the message text.
8. In the buttons, icons, and explanation and responses sections, specify additional details.
9. Click **New Row** to add more messages, or save your changes.

Query definitions

Users typically have a well-defined set of columns that they want to query in each application. You can identify these columns by asking users or by examining reports of slow-running SELECT statements. Indexing these columns can improve system performance.

Users can create and save their own queries and share queries with other users. Saved queries are stored in a table named QUERY. You must periodically review these saved queries for inefficient conditions and use of unindexed columns.

With SQL statements, you can create special-purpose queries for example, return all Preventive Maintenance work orders created since Monday of this week. When you create these queries, you can save users the effort of querying larger sets and sorting and scrolling through the sets.

You also can provide users with an efficient default query for frequently used applications so that they can see a preferred record set when they access the application. For example, in the Work Order Tracking application, you can specify a default query for supervisor Smith. With this default query in place, initially the user can see only work orders with SMITH in the **Supervisor** field.

Search option configuration for performance optimization

You want to prevent users from performing queries that retrieve hundreds of thousands of records. You should periodically review users saved queries, you could define default queries for your users, or educate them on how to create efficient queries

Text search function

The text search function uses SQL to index, search, and analyze text and documents that are stored in system databases. You can search for information based on text, content metadata, or attributes.

A text search is allowed only for fields of the ALN data type. A text search is designed to search long descriptions, or fields that are long data types.

A full text search is a language-specific text search. Parts of words, not words, are indexed. For example, if you search for *par*, results do not include *part*.

A text search is flagged on the object and on the attributes.

Search type configuration

By setting search types for database columns, you can improve the results that are returned by user queries. Setting the search type can also reduce the load on the database.

Tables with fewer than 2000 or 3000 records are typically scanned regardless of indexes. The input output (I/O) cost to read an entire table is less than the average I/O cost of the index lookup plus the table lookup. The SEARCHTYPE value does not affect database behavior when such scans are performed. Tables with relatively few rows have no noticeable degradation in performance.

The following search types are available for user queries.

Table 18. Search types for user queries

Search type	Description	Benefit	Entry required
Exact	Filters data based on the keywords that you specify. You specify an exact search type when you require accurate and targeted results. Exact searches use wild cards only if a user explicitly enters wildcard characters on the List tab or in the WHERE clause.	Key fields, such as Work Order and Purchase Order, and value list fields, such as Work Order Status, can benefit from the indexing that is used in exact searches.	=
Wildcard	The default search type is wildcard search. You can apply a wildcard search on description fields of tables that have a relatively small number of rows for example, 2000 or fewer rows. When a user enters a value in a field on the List tab, the wildcard search type condition looks like this: column like '%value%' In wildcard searching, the database engine cannot use indexes. Searching without indexes can result in slower search times, especially on tables with many rows.	Wildcard searching provides flexibility for the users.	%

Table 18. Search types for user queries (continued)

Search type	Description	Benefit	Entry required
Full text	You can specify a text search type on description fields of tables with large numbers of rows, for example, tens of thousands of rows. The text search engine takes time to refresh the indexes, so new records might not be found until the text search index refreshes itself. Stem search is also performed. For example, a search for service returns servicing and serviced. Most system tables have one or more ALN data type columns for descriptions, memos, or remarks. You can define text search types and a corresponding Oracle text index or SQL Server full text catalog for columns that have excessive text. • On Oracle, you can modify the <code>maximo_ts_job_call</code> procedure to change the schedule of the synchronization process to any interval. • On SQL Server, you can set and modify the population schedule for the full text catalog. • Full text indexing is not available on IBMDB2. Text indexing increases the load on the database because of the constant background processing to keep the indexes synchronized. However, text indexing produces efficient word searching of description fields.	Text searches produce faster search responses than wildcard searches. Fields that are text-search enabled have text search indexes and, therefore, result in a faster search response. If an object is enabled for text search, the full text searches on the attributes provided.	Any combination of the words in the text search
None	The none search type is used for columns that cannot be searched. If you do not specify a search type, the value defaults to none or no search.	You use this search type to specify that a column should not be searched.	

You can use a combination of methods to refine searches.

Application Designer application

You can use the Application Designer application to customize an application by adding or removing columns from the **List** tab. You can then ensure that the columns to be queried are all indexed.

Application cloning feature

You can clone an application and then use the Application Designer application to create an alternate version with a restricted number of columns that can be queried.

Security groups feature

After you clone applications, you can use security groups to assign users to specific application clones. You can also use security groups to prohibit access to the **More Search** fields and WHERE clause advanced query options. When you prohibit access to those options, you limit users to query on the **List** tab of the application.

Chapter 3. Configuring the system with multiple languages

A database can contain data in multiple languages. Diverse users can run the system in their native languages. By default, multiple languages are enabled for data dictionary tables, company and item objects, and system messages. You use the translation data toolkit (TD Toolkit) to access non-English language databases.

Configuration of multiple languages overview

You can configure the database to meet the needs of your users. Translate a database in multiple languages using the translation data toolkit (TD Toolkit.)

Multiple language tables and associated columns

Enabling multiple languages on an object or on a table creates a secondary table connection. For example, L_ITEM is the secondary table for the ITEM object.

You use the MAXATTRIBUTE table to determine which tables and columns are multiple language supported or multiple language enabled. The flag for multiple language supported (MLSUPPORTED=1) indicates whether the values in the column are stored in multiple languages or not. This flag is read-only and cannot be changed.

The flag for multiple language enabled (MLINUSE=1) indicates whether the column is enabled to store the values in multiple languages.

To view the tables and columns that are enabled for multiple languages, access a Structured Query Language (SQL) editor and type `select objectname,attributename from maxattribute where mlinuse= 1;`

By default, the following tables and associated columns are enabled for multiple languages.

Table 19. Multiple language tables and associated columns

Tables	Columns
ALNDOMAIN	DESCRIPTION
ASSETATTRIBUTE	DESCRIPTION
COMMTEMPLATE	MESSAGE SUBJECT
COMPANIES	NAME NAME_LONGDESCRIPTION
CTRLCONDPROP	CTRLCONDPROP
ITEM	DESCRIPTION_LONGDESCRIPTION DESCRIPTION
MAXAPPS	DESCRIPTION
MAXATTRIBUTE	TITLE REMARKS

Table 19. Multiple language tables and associated columns (continued)

Tables	Columns
MAXATTRIBUTECFG	TITLE REMARKS
MAXDOMAIN	DESCRIPTION
MAXLABELS	VALUE
MAXMENU	HEADERDESCRIPTION
MAXMESSAGES	EXPLANATION ADMINRESPONSE BUTTONTEXT SYSTEMACTION VALUE OPERATORRESPONSE
MAXMODULES	DESCRIPTION
MAXOBJECT	DESCRIPTION
MAXOBJECTCFG	DESCRIPTION
MAXSERVICE	DESCRIPTION
NUMERICDOMAIN	DESCRIPTION
PALETTEITEM	DESCRIPTION
REPORT	DESCRIPTION
REPORTLABEL	LABELVALUE FONTNAME
REPORTLOOKUP	LABELOVERRIDE
SIGOPTION	DESCRIPTION
SOLUTION	DESCRIPTION
SYNONYMDOMAIN	DESCRIPTION

Related tasks:

“Enabling multiple languages on objects and attributes” on page 88

You can enable multiple languages on objects or attributes.

“Enabling attributes for multiple languages” on page 88

You can enable attributes for multiple languages.

“Viewing characters from multiple languages” on page 89

To view characters from multiple languages, you can install additional language files.

Multiple language utilities - translation data toolkit

To use non-English-language databases, you use the translation data toolkit (TD Toolkit). The TD Toolkit is a utility that extracts translatable data from a database.

You can use the data as a translation of the database or add the data as a secondary language. Whether the secondary language is supported determines how you can use the data.

Maximo Asset Management is translated into 24 languages and the TDToolkit supports all 24 languages

You also can use the TD Toolkit to translate the report labels and parameter labels for reports.

Related tasks:

“Adding unsupported second languages to databases” on page 90

You use the Translation Data toolkit (TD toolkit) to add unsupported language to databases. The base language remains the same, and the other language is added as a second language. For example, a Canadian company can add French as a second database language.

Multiple languages and system table customizations

When you use a multiple language implementation, you can customize system tables to fit your business needs.

The following table lists some of the system tables that you can customize.

Table 20. Examples of system tables that can be customized

Table	Description
MAXATTRIBUTE	Stores information that is associated with individual object attributes
MAXLABELS	Stores application labels that are associated with individual application fields
MAXMENU	Stores menu values associated with individual applications
MAXMESSAGES	Stores application messages that are associated with popup boxes and buttons

Multiple languages and translations

If you are using a multiple language implementation, you can track and perform translations on new records. By default, new records are stored in the base language only; there is no auto-translation.

You can translate records to the secondary language using one of the following options:

- Translate each record in the localized application. Individual record translations occur only on implementations that require minimal translations.
- Translate your records through the resulting XLIFF file from the translation data toolkit utility, `TDToolkit.bat`.
- Translate your records through the resulting XLIFF file from the translation data toolkit utility, `TDToolkit.bat`. Follow the procedure to add unsupported second languages to database, begin at step 3 because you already translated the XLIFF files. Add the translated XLIFF files into the database.
- Translate your records through the resultant XML file from the `exportlang.bat` utility.

XML Localization Interchange File Format (XLIFF) is an XML file with specific tags for the translation process. XLIFF files follow the 8-bit Unicode Transformation Format (UTF-8) standard for formatting languages.

Related tasks:

“Localizing databases for unsupported base languages” on page 89

You use the translation data toolkit (TD toolkit) to localize the database for an unsupported base language.

“Adding unsupported second languages to databases” on page 90

You use the Translation Data toolkit (TD toolkit) to add unsupported language to databases. The base language remains the same, and the other language is added as a second language. For example, a Canadian company can add French as a second database language.

Enabling multiple language support

You can enable multiple languages on objects or attributes. You can use the translation data toolkit (TD toolkit) to localize the database for unsupported base languages. You can also use the TD toolkit to add supported languages to databases.

Enabling multiple languages on objects and attributes

You can enable multiple languages on objects or attributes.

Procedure

1. In the Database Configuration application, select the object (for example, ASSET or LOCATIONS) that you want to enable for multiple languages.
2. In the **Objects** tab, specify a value for the Language Table. The convention is *L_objectname*.
3. Save the record.

Related concepts:

“Multiple language tables and associated columns” on page 85

Enabling multiple languages on an object or on a table creates a secondary table connection. For example, L_ITEM is the secondary table for the ITEM object.

Enabling attributes for multiple languages

You can enable attributes for multiple languages.

About this task

You use the Database Configuration application to create language objects.

Most of the system attributes do not support multiple languages. For example, description fields in ITEM and COMPANIES support multiple languages, while description fields in transaction applications, such as WO, PO, PR, RFQ, and INVOICE, do not. The language tables are empty until you populate them with data. You can export and import all translatable strings through XLIFF files.

Procedure

1. In the Database Configuration application, select the attribute that you want to enable for multiple languages.
2. From the **Attributes** tab, check that **Multilanguage Supported** is selected.
3. Select the **Multilanguage in Use** check box to identify the attributes that you want to enable for multiple languages.
4. Configure the database.

Related concepts:

“Multiple language tables and associated columns” on page 85

Enabling multiple languages on an object or on a table creates a secondary table connection. For example, L_ITEM is the secondary table for the ITEM object.

Viewing characters from multiple languages

To view characters from multiple languages, you can install additional language files.

Before you begin

The files for East Asian languages (Chinese, Japanese, and Korean) require 230 MB of disk space. The files for complex script and right-to-left languages (Arabic, Armenian, Georgian, Hebrew, Indic languages, Thai, and Vietnamese) require 10 MB of disk space.

About this task

Certain fonts do not support foreign language characters. For example, Veranda does not support East Asian characters.

Procedure

1. From the **Start > Control Panel** menu, select **Regional and Language Options**.
2. On the **Languages** tab, select one of the following check boxes:
 - **Install files for complex script and right-to-left languages (including Thai)**
 - **Install files for East Asian languages**
3. Click **OK** or **Apply**, and restart your computer.

Related concepts:

“Multiple language tables and associated columns” on page 85

Enabling multiple languages on an object or on a table creates a secondary table connection. For example, L_ITEM is the secondary table for the ITEM object.

Localizing databases for unsupported base languages

You use the translation data toolkit (TD toolkit) to localize the database for an unsupported base language.

Before you begin

XML Localization Interchange File Format (XLIFF) is an XML file with specific tags for the translation process. XLIFF files follow the 8-bit Unicode Transformation Format (UTF-8) standard for formatting languages.

Procedure

1. To export translatable data from the database, complete the following steps:
 - a. Change to the `tools\maximo` installation subdirectory and run the export command.
 - b. Run the `tdtoolkit -export batch file`. This batch file exports translatable data into the `\tools\maximo\xliff\export` directory.
2. To translate all XLIFF files from English to the localized language, complete the following steps:
 - a. In each XLIFF file, find the instances of the `<target> </target>` tags.

- b. Replace the text between the tags with the text in the language into which you are translating.
3. To import the translated XLIFF files into the database, complete the following steps:
 - a. Create the subdirectory, `\tools\maximo\lc\xliff` in which *lc* represents the two-letter language code. The two-letter language code is in the MAXLANGCODE column in the LANGUAGE table.
 - b. Copy the files in the `\tools\maximo\xliff\export` directory to the new subdirectory, `\tools\maximo\lc\xliff`.
 - c. To translate the XLIFF file into the database, in the installation directory, run the `tdtoolkit -import -tlc -versionV7100-000 import` command. For example, to import a Hebrew (unsupported) database, run the import command: `tdtoolkit -import -tlhe -version7100-000`.

Example

If the database is used in another base language, you can translate the English language database into an unsupported base language. For example, you can translate a database to Arabic for Arabic-speaking users. You can add a secondary language, such as French, to the Arabic language database.

Related concepts:

“Multiple languages and translations” on page 87

If you are using a multiple language implementation, you can track and perform translations on new records. By default, new records are stored in the base language only; there is no auto-translation.

Adding unsupported second languages to databases

You use the Translation Data toolkit (TD toolkit) to add unsupported language to databases. The base language remains the same, and the other language is added as a second language. For example, a Canadian company can add French as a second database language.

Before you begin

XML Localization Interchange File Format (XLIFF) is an XML file with specific tags for the translation process. XLIFF files follow the 8-bit Unicode Transformation Format (UTF-8) standard for formatting languages.

About this task

You can add an unsupported second language to the database to allow non-English language users to view and make database changes. If the database is used in a base language other than English, you can translate the English language database into an unsupported base language. For example, you can translate a database to Estonian for Estonian-speaking users. You can add a secondary language, such as French, to the Estonian language database.

Procedure

1. To export translatable data from the database, complete the following steps:
 - a. In the `tools\maximo` installation subdirectory, run the export command.
 - b. Run the `tdtoolkit -export batch file`. This batch file exports translatable data into the `\tools\maximo\xliff\export` directory.
2. To translate all XLIFF files from English to the localized language, complete the following steps:

- a. In each XLIFF file, find every instance of the <target> </target> tags.
 - b. Replace the text between the tags with the text in the language into which you are translating.
3. To add the translated XLIFF files into the database, complete the following steps:
 - a. Create the \tools\maximo\lc\xliff subdirectory in which *lc* represents the two-letter language code. The two-letter language code is in the MAXLANGCODE column in the LANGUAGE table.
 - b. Copy the files in the \tools\maximo\xliff\export directory to the new subdirectory, \tools\maximo\lc\xliff.
 - c. To translate the XLIFF file into the database, in the installation directory, run the `tdtoolkit -import -tl1c -versionV7100-000 import` command. For example, to import an Indonesian (unsupported) database, run the import command: `tdtoolkit -import -tlid -version7100-000`.

Related concepts:

“Multiple language utilities - translation data toolkit” on page 86

To use non-English-language databases, you use the translation data toolkit (TD Toolkit). The TD Toolkit is a utility that extracts translatable data from a database.

“Multiple languages and translations” on page 87

If you are using a multiple language implementation, you can track and perform translations on new records. By default, new records are stored in the base language only; there is no auto-translation.

Removing secondary languages from the database

If you have secondary languages installed on your database that you no longer need, you can remove them. Removing secondary languages frees memory and might improve the performance of the database.

Procedure

1. Optional: Run the following SQL statement to get a list of all languages installed on the database including the base language: `select languagename, maxlangcode from language where enabled = 1;`
2. Optional: If you do not know the base language, run the following SQL statement: `select varvalue from maxvars where varname = 'BASELANGUAGE';` When the base language is identified, you can identify secondary languages.
3. To disable a language, run the following SQL statement on the database, replacing *language_code* with the language code of the secondary language that you want to remove: `update language set enabled = 0 where maxlangcode = 'language_code';` This command does not remove the translation of the secondary language. Ensure that you do not disable the base language.
4. To remove the translation of the secondary language, run the following SQL statement against your database, where *language_code*, is the language code of the language to remove:

Database	SQL Statement
Oracle or IBM DB2	<code>select 'delete from ' langtablename ' where langcode = ' language_code ';' from maxtable where langtablename is not null;</code>
Microsoft SQL Server	<code>select 'delete from ' + lower(langtablename) + ' where langcode = ' language_code ';' from maxtable where langtablename is not null;</code>

The output from this statement is approximately 32 lines of code.

5. Copy and paste all the delete statements from the output and run the output against your database.
6. To remove the language from the table that is used to track versions, run the following SQL statement against your database:

```
delete from tdtversion where language = 'language_code';
```

What to do next

You must restart the application server after you remove a language for the change to take effect.

Translating records through applications

If you are using multiple languages, you can translate records to the secondary language through an application.

Procedure

1. Select the appropriate secondary language and log in.
2. Access the application that contains the records that you want to translate.
3. Select the records that you want to translate. For example, in the Item Master application, search for and select the item numbers that you want to translate.
4. Save the records. The records are saved to the secondary language ITEM table.

Setting languages for tooltips

You can change the language for tooltips in the ilog toolbar of the applet title bar in the Assets application. You must reconfigure the Java virtual machine (JVM) for the applet or the application on which the browser is running for the change to take effect.

About this task

This ilog toolbar includes the following tools:

- **Make Select Active**
- **Zoom In**
- **Zoom Out**
- **Zoom Box**
- **Reset Zoom**
- **Fit to Contents**
- **Pan**

Procedure

1. On your desktop, click **Start > Control Panel**.
2. Click the Java program and click the **Java** tab.
3. Select the **View** button for the setting the browser uses:
 - Java Applet Runtime
 - Java Application Runtime
4. In the table, click row for the JVM your browser is using.

5. Double click the **Runtime Parameters** cell and type
`-Duser.language=language_code` and `-Duser.country=language_code` For example, enter:
`-Duser.language=ru -Duser.country=RU`
6. Select the **Enable** check box.
7. Click **Apply** and click **OK** .

What to do next

You must restart the browser for the changes to take effect.

Create a maxdemo database after installation

You can create a maxdemo database and additional databases after you have installed the product. You can populate the database with sample data. You use the maxinst program to create additional databases.

Before you begin

The maxinst program does not provide default values for table space parameters. Specify the data and index table space names to ensure that your installation runs smoothly.

About this task

If you created a database either automatically or manually during the installation, you can use maxinst to optionally create a maxdemo database in that database instance. If the maxinst program fails, you must recreate the Maximo database schema before running the maxinst program again.

Procedure

1. Open a command window and change directory to `C:\IBM\SMP\Maximo\tools\maximo`.
2. You can create an additional database in one of the following ways:
 - To create a DB2 database, run the following command to set the DB2DBDFT variable:
`set DB2DBDFT=dbname`
 - To create an SQL Server database,
 Launch maxinst using the `-s` and `-t` parameters:
`maxinst -sPRIMARY -tPRIMARY -m2`
 - To create an empty Maximo database for Oracle, run the following command:
`maxinst -s<tablespacename> -t<tablespacename> -imaximo`
 - To create a maxdemo database for Oracle, run the following command:
`maxinst -s<tablespacename> -t<tablespacename>`

For example, type `maxinst -sMAXIMO -tMAXIMO`.

The system reads the maximo.properties file for database connectivity information. The maximo.properties file is located in the `C:\IBM\SMP\Maximo\Applications\Maximo\Properties` directory. The system connects to the database through the JDBC connection and creates a maxdemo database.

3. You can populate the additional database by running commands with specific parameter values. The following table lists the maxinst database parameters:

Parameter	Description
-a	Database alias. If not specified, the alias <code>mxe.db.url</code> property is used.
-d	Log file directory. If you are using the <code>-l</code> parameter, the log file is sent to the specified directory. Otherwise, the log file is sent to the log directory, for example <code>C:\IBM\SMP\Maximo\tools\maximo\lo</code> .
-e	Runs the SQL. This parameter is required and already present in the <code>maxinst.bat</code> file.
-f	File name for the properties file. If not specified, <code>maximo.properties</code> is used.
-i	File name of the input file (without path or extension). If not specified, the default file name <code>Unlcvt</code> is used.
-k	Directory of the properties file.
-l	Creates a detailed log file. This parameter is already present in the <code>maxinst.bat</code> file.
-o	If you are using the <code>-l</code> parameter, the <code>-o</code> parameter specifies the file name for the log file.
-p	Password for the database connection. If not specified, the <code>mxe.db.password</code> property or <code>MAXIMO</code> is used. If <code>MAXIMO</code> is used, it must be entered in uppercase letters.
-s	Required: Table space for index storage.
-t	Required: Table space for table storage.
-u	User name for database connection. If not specified, the <code>mxe.db.user</code> property or <code>MAXIMO</code> is used. If <code>MAXIMO</code> is used, it must be entered in uppercase letters.
-x	Required for UNIX: Fixes the doclink file separators in UNIX environments. Note: If a UNIX environment is deployed without using this parameter, the attached documents do not function properly.

4. Add the installation-related properties to the database from the `install.properties` file. The `install.properties` file is located in the `C:\IBM\SMP\ETC` folder. You can add these properties to the database from the System Properties application.

Chapter 4. Administering the database

To maintain database integrity, you must perform backups and other tasks on a regular basis. Before you perform a restoration procedure, test the process in a test environment, even if your backup procedure appears to be working properly. See your database platform documentation for specific commands and procedures to restore your database from a backup.

Database administration overview

You can schedule system and database backups to run on a regular basis to ensure that you have a recent backup of all your information. After you apply a patch or install new system options, you can use the database update utility to synchronize the system and the database.

Database backup and restoration

Backup procedures depend on the size of your database and the type of operation you are running.

You might want to store backups in a different location from your production database and application files. You also might want to schedule and regularly perform system and database backups.

You can back up using various archive mediums. The following table lists the types of archive media that you can use.

Table 21. Types of archive media

Media	Description
Hard disk drive	Lets you restore your system quickly.
Tape drive	This method is slower, but you can keep multiple tapes of backups. Tape drive backups typically include backup software; see the drive software documentation.
CDs, DVDs, diskettes	These devices offer limited capacity. However, the devices are useful for smaller databases, for archive files, or for specific executables.

Types of backups

There are different ways to back up your system. These methods vary in scope and frequency. See your database platform documentation for specific commands and procedures to perform backups.

The table that follows lists and describes the types of backups you can perform.

Table 22. Types of backups

Type	Description	Frequency
System backup	Completely duplicates the system software. Lets you restore the entire system to its original state, including customized applications and reports. Backs up all system product files (on the administration workstation) as well as application server product files (on the application server). On LAN systems, perform system backups when all users are logged out of the system.	As needed, when you modify software or reports.
Database backup	Duplicates only the databases.	Daily to ensure full recovery of data no more than a day old. After long data entry sessions. At the end of accounting and reporting periods. Before any critical event, such as an outage or plant turnaround. Before and after configuring the database. Before and after installing patches and product add-ons.

Offline and online backups

Offline and online backups allow you to back up your system whether the system is being actively used or not. See your database platform documentation for specific commands and procedures to perform backups.

The following table lists the types of database backups that you can perform.

Table 23. Types of database backups

Type	Description
Offline backups (standard)	Perform offline database backups with all users logged out of the system and the database server down. Duplicates of the database made while the server is up and users are connected can result in unrecoverable backups. Shut down the application server and report server before you perform the backup. When the backup is complete, restart the database server, application server, and report server.
Online backups	You can perform backups without bringing the database server down, which lets the users continue to use the software during the backup. This process is more time-consuming, but can be useful to minimize downtime in 24-hour operations.

Database statistics updates

To enhance performance, regularly update your database statistics. See your database platform documentation for procedures.

DBMS_STATS package

The DBMS_STATS package in Oracle optimizes statistics on your database. The system benefits from cost-based optimization because it builds many queries dynamically, depending on user input.

With the cost-based optimizer, Oracle determines which indexes to use based on the distribution of data. Oracle 9i and 10 g documentation recommend against using ANALYZE to collect statistics for the Cost Based Optimizer; use DBMS_STATS instead.

Update statistics

If your database is large, run the Oracle update statistics. You can use a database-specific command, or you can run update statistics from the actions menu in the Database Configuration application. This calls dbms_stats.gather_table_stats with cascade true.

Update statistics example

```
dbms_stats.gather_table_stats
(ownname => 'MAXIMO', tabname => 'ASSET', cascade => true)
```

Optimizer modes

Oracle has two optimizer modes: cost-based and rule-based.

By default, the optimizer mode is set to CHOOSE. To determine the mode in effect, select from the v\$parameter table:

```
select value from v$parameter where name='optimizer_mode';
```

If the mode is CHOOSE, you use the rule-based optimizer unless statistics exist. Statistics do not exist if you never analyzed your tables. Do not set the optimizer mode to RULE.

SQL server update statistics

Perform the update statistics procedure to ensure that selectivity factors are updated when there are significant changes to an index.

You might want to perform this procedure daily, especially if large amounts of data are inserted, updated, or deleted. You can run update statistics from the actions menu in the Database Configuration application, or use a database-specific command.

Database updates

The system includes a Maximo database update utility called UpdateDB.

Run UpdateDB after you install system application patches or after you install any system options. System options might include Maximo, IBM Maximo Mobile, or other applications. After you install patches or options, your application version will be different from your Maximo database version. For the system to function properly, the system and Maximo database versions must match.

When you start the application server (MXServer) the system compares the application version to the Maximo database version. If the system detects a discrepancy, the MXServer stops processing and the system prompts you to run the system UpdateDB utility. The upgrade script and class files run during the database update and revise the version references in the Maximo database, synchronizing the system, and database versions.

Application patches

Application patches are available for download on the IBM Software Support site.

Database update for system options

All system options use the `a_customer.xml` file and the `product_description.xml` files in the update process. These files are located in the `maximo\properties\product` folder.

The .xml files listed contain the following information:

- Dbmaxvarname – database maxvar name for the system option
- Dbscripts – script directory name where system product script files are located
- Dbversion – current system option version
- Lastdbversion – last release version
- Extensions – class file extension information for system option

The first file run by the UpdateDB utility is the `a_customer.xml` file. Next, the update utility runs each of the `product_description.xml` files in alphabetic order.

The UpdateDB utility is configured to run scripts based on the values specified in each of your .xml files. The scripts representing each successive update version up to and including the referenced dbversion value script are run during your database update process. Upon completion, your dbversion value is updated to the most current script version value.

UpdateDB and customer extensions

When you run `updatedb.bat`, you receive a message for certain customer extensions.

Product {Industry solution name} has extensions but `a_customer.xml` file does not exist. Do you want to continue (Y/N)?

- If you type Y, the UpdateDB process continues.
- If you type N, the UpdateDB process stops.

a_customer.xml file

The system uses the a_customer.xml file to reference any system classes that have been customized.

Because the a_customer.xml file is the first to be run by the UpdateDB utility, the changes you reference in the product script files are the first to be applied. All your system options are then incorporated into the customization before the UpdateDB utility runs the product_description.xml scripts. If you incorporate class extensions in any of your system options, create the a_customer.xml file. All modified class files and scripts must be referenced in the format shown in the following example.

Example

In this example, the UpdateDB utility runs the scripts representing each successive update version up to and including the referenced V600_01 script. The altered <mboaset objectname> and <mbo object> entries indicate that the purchase order classes have been extended.

```
<?xml version="1.0" encoding="UTF-8"?>
<product>
  <name>Customer Product</name>
  <version>
    <major>6</major>
    <minor>0</minor>
    <patch>0</patch>
    <build>999</build>
  </version>
  <dbmaxvarname>DBCUST</dbmaxvarname>
  <dbscripts>cus</dbscripts>
  <dbversion>V600-01</dbversion>
  <lastdbversion>V520-20</lastdbversion>
  <extensions>
    <mboaset objectname='P0'>psdi.app.cust.P0Set</mboaset>
    <mbo objectname='P0'> psdi.app.cust.P0</mbo>
  </extensions>
</product>
```

product_description.xml file

The product_description.xml file identifies each system option installed on your system.

For each industry solution that you installed, create a separate <productname>.xml file to deploy the EAR files successfully. Create new <productname>.xml files in the new maximo\properties\product directory.

```
<?xml version="1.0" encoding="UTF-8"?>
<product>
  <name>IT and Service Management Extensions</name>
  <version>
    <major>6</major>
    <minor>0</minor>
    <patch>0</patch>
    <build>999</build>
  </version>
  <dbmaxvarname>DBITSME</dbmaxvarname>
  <dbscripts>itsme</dbscripts>
  <dbversion>V600-01</dbversion>
  <lastdbversion>V520-20</lastdbversion>
  <depends>newproduct</depends>
```

```

<extensions>
  <mboaset objectname='objectname'>classname</mboaset>
  <mbo objectname='objectname'>classname</mbo>
  <field objectname='objectname'
    attributename='attrname'>classname</field>
  <service servicename='servicename'>classname</service>
  <bean presentation='appname' controlid='id'>beanclassname</bean>
  <class extends='classname'>classname</class>
</extensions>
</product>

```

Managing database administration

You can ensure that your database is working at optimum performance levels by applying patches and running regular updates.

Updating the Maximo database

In order to keep it up to date, you must periodically apply patches to the Maximo database.

Procedure

1. Download and apply the application patch.
2. Back up the database.
3. Run updatedb.bat to update the database.

Related concepts:

“Database backup and restoration” on page 95

Backup procedures depend on the size of your database and the type of operation you are running.

“Types of backups” on page 95

There are different ways to back up your system. These methods vary in scope and frequency. See your database platform documentation for specific commands and procedures to perform backups.

“Offline and online backups” on page 96

Offline and online backups allow you to back up your system whether the system is being actively used or not. See your database platform documentation for specific commands and procedures to perform backups.

“Database statistics updates” on page 97

To enhance performance, regularly update your database statistics. See your database platform documentation for procedures.

“Database updates” on page 98

The system includes a Maximo database update utility called UpdateDB.

Running UpdateDB

Class files for running UpdateDB are located in the *install_home\tools\maximo\classes\psdi\script\en* directory. Script files are located in the *install_home\tools\maximo\en* directory.

Procedure

1. Open a command prompt and change the directory to: *install_home\tools\maximo*.
2. At the prompt, type *updatedb.bat* and press **Enter**.

What to do next

If you encounter problems during the system update process, the system logs errors to the *install_home\tools\maximo\logs\Update+Timestamp.log* file. You can examine the logs to determine the source of update errors.

When you successfully complete a database patch update, the database build version in the MAXVARS table is revised.

Related concepts:

“Database backup and restoration” on page 95

Backup procedures depend on the size of your database and the type of operation you are running.

“Types of backups” on page 95

There are different ways to back up your system. These methods vary in scope and frequency. See your database platform documentation for specific commands and procedures to perform backups.

“Offline and online backups” on page 96

Offline and online backups allow you to back up your system whether the system is being actively used or not. See your database platform documentation for specific commands and procedures to perform backups.

“Database statistics updates” on page 97

To enhance performance, regularly update your database statistics. See your database platform documentation for procedures.

“Database updates” on page 98

The system includes a Maximo database update utility called UpdateDB.

Chapter 5. Optimizing system performance

System performance is dependent on all aspects of your deployment. You can tune the database, application servers, and HTTP servers, and you can also improve the performance of the operating system. Performance tests can provide you with results that you might use to further adjust configuration settings for optimal performance.

Related information:

Implementing high availability

Performance improvements with clusters

A cluster groups similar functions on two or more Java virtual machines (JVMs) to process a single function, such as scheduled cron tasks. Clusters connect to the same database but operate independently. For example, if the cron task cluster fails, users can still connect to the user interface cluster.



Performance Best Practices White Papers (Opens in a new browser window or tab.)



Maximo Activity Dashboard (PerfMon) (Opens in a new browser window or tab.)



Interactive Performance Analysis with Maximo Activity Dashboard (Opens in a new browser window or tab.)



Tuning WebLogic Server (Opens in a new browser window or tab.)

Database server performance

The database server stores all data that is collected and calculated by Maximo Asset Management applications. The database server also stores configuration metadata for your system environment, processes all transactions from the applications, and handles requests for management reports. A properly configured and maintained database is required for performance optimization.

Related tasks:

“Developing performance tests” on page 130

Performance tests provide baseline measurements of system performance, which you can use to identify the weaknesses in your deployment and to determine optimal configuration settings. When you design performance tests, you determine test objectives, develop use cases, develop a test strategy, and define your test environment.

Related information:

High availability for databases

Optimization techniques for all databases

You can apply standard database-tuning techniques to Maximo Asset Management by periodically monitoring a production database during peak load. You can use any appropriate monitoring tools or techniques and, if necessary, adjust parameters to resolve the reported problems.

Related tasks:

“Optimizing performance in DB2” on page 107

You can change configuration settings at the registry, database manager, and the database levels. The registry level controls all DB2 databases and DB2 applications. The database manager controls the main configuration settings for all databases. Each individual database instance can also have its own settings at the database level.

“Optimizing performance in Oracle Database” on page 117

Initialization parameters set values that can affect system performance, such as the optimizer features. Initialization parameters are stored in an initialization parameter file. You can change initialization parameters with **ALTER SYSTEM** commands.

Related information:

 Improving the performance of the doclinks query in process automation engine products (Opens in a new browser window or tab.)

 Keeping your Maximo Tuned Up and Performing (Opens in a new browser window or tab.)

 Working with Archived Data (Opens in a new browser window or tab.)

High availability for databases

Configuring databases for high availability

Database indexing

Indexing a database requires good understanding of the data, the user functions, and how the database is indexed. Indexes use key parts of data from a table in a binary structure to improve searching capability. Each record of data in the table must be associated with data in the index.

Indexing can greatly increase search speeds. However, a drawback of indexes is that each insert, update, or delete operation requires an update to the indexes. When tables include multiple indexes, each index can increase the time it takes to process table updates. If you want to reduce the number of indexes to improve processing speed, remove indexes that are least valuable for search purposes.

Some index types that are available on DB2, Oracle Database, and Microsoft SQL Server are not available in the Database Configuration application. You can create and maintain these indexes from the command line to improve performance in specific cases. For example, on Oracle Database, you can create a bitmap index or a function-based index if you determine that these indexes can improve certain queries.

If you use these index types, the system administrator must remove any of these indexes before you configure the database changes. After the database is configured, the indexes must be replaced.

If you customize Maximo Asset Management, you can change the way you select information from the database. For example, a customization might include additional tables and columns. If you customize Maximo Asset Management, compare indexes to the user functions that use them. Ensure that you implement the right balance of indexes.

Related concepts:

“Reorganization of tables and indexes in DB2” on page 116

When many updates are made on a table, the space can become fragmented. You can reorganize data to reclaim space in a table and to improve system performance. Use the IBM DB2 **REORGCHK** and **REORG** commands to optimize table spaces and indexes. You cannot reorganize metadata tables.

Optimized access to data

The use of sequence cache and separate table spaces for large tables can help you improve your system performance. Also, you can set applicable page sizes and storage capacities to these table spaces and cache memories for optimal performance.

Table spaces

A database administrator can use the following guidelines for table space page sizes. The administrator then moves the tables that include the most record types and are the most heavily used into table spaces with these page sizes:

Table 24. Guidelines for table spaces page sizes

Page size	Row size	Column count limit	Maximum capacity (DMS table space)
4 KB	4005	500	64 GB
8 KB	8101	1012	128 GB
16 KB	16293	1012	256 GB
32 KB	32677	1012	512 GB

Separate indexes from the data and place them into a different table space. Use the Database Configuration application to move the indexes. Also, separate large tables, such as Assets and Work Orders, into their own table spaces for optimal performance.

Sequence cache

Sequence caches are used to automatically generate unique values in the database. These values are typically used as identifiers for columns or rows. Set the maxseq sequence for the rowstamps column to 500 because this column frequently requires values. Set the cache size for all other sequences to 50.

Related tasks:

“Modifying sizes of sequence caches”

A sequence is a database object that automatically generates unique key values. A sequence can generate one value at a time or can generate a cache of multiple values. The use of sequence caches improves system performance because processes can obtain values from the cache without waiting for the sequence to generate individual values.

Modifying sizes of sequence caches

A sequence is a database object that automatically generates unique key values. A sequence can generate one value at a time or can generate a cache of multiple values. The use of sequence caches improves system performance because processes can obtain values from the cache without waiting for the sequence to generate individual values.

About this task

The maxseq sequence is used to generate values for the rowstamps column in database tables. Because the rowstamps column requires values frequently, set the sequence cache size for maxseq to 500. For all other sequences, set the sequence cache size to 50.

If IBM Tivoli® Asset Management for IT is installed, you must remove some specified sequence names from your SQL script because those sequences must not be altered manually.

Procedure

1. Run the following command to generate a script file that contains the SQL statements to set the sequence cache size:

Option	Description
If you are using DB2	db2 "select 'alter sequence maximo.' sequencename ' cache 50 ;' from maximo.maxsequence" > change_seq_cache.sql
If you are using Oracle	sqlplus "select 'alter sequence maximo.' sequencename ' cache 50 ;' from maximo.maxsequence" > change_seq_cache.sql

2. Edit the change_seq_cache.sql file to change the sequence cache size for the maxseq sequence to 500.
3. If Tivoli Asset Management for IT is installed, edit the change_seq_cache.sql file to remove any entries that match the following sequence names:

Table 25. Sequences to remove from the change_seq_cache.sql file

Sequence cache name	Sequence cache name	Sequence cache name
ASSETATTRIBUTESEQ	DPADISKSEQ	DPAMSWSUITECOMPSEQ
CDMCITYPESSEQ	DPADISPLAYSEQ	DPAMSWSUITESEQ
CLASSANCESTORUSEQ	DPADFILESEQ	DPAMSWUSAGERANGESEQ
CALSSSPECSEQ	DPAIMAGEDEVICESEQ	DPAMSWUSAGESEQ
CLASSSPECUSEWITHSEQ	DPAIPXSEQ	DPAMSWVARIANTSEQ
CLASSSTRUCTURESEQ	DPALOGICALDRIVESEQ	DPANETADAPTERSEQ
OMPSEQ	DPAMADAPTERSEQ	DPANETDEVCARDSEQ
RELATIONRULESEQ	DPAMADPTVARIANTSEQ	DPANETDEVICESEQ
RELATIONSEQ	DPAMEDIAADAPTERSEQ	DPANETPRINTERSEQ
ACTCIRELATIONSEQ	DPAMMANUFACTURERSEQ	DPAOSSEQ
ACTCISEQ	DPAMMANUVARIANTSEQ	DPASOFTWARESEQ
ACTCISPECSEQ	DPAMOSSEQ	DPASWSUITESEQ
DEPLOYEDASSETSEQ	DPAMOSVARIANTSEQ	DPATCPIPSEQ
DPACOMMDEVICESEQ	DPAMPROCESSORSEQ	DPAUSERINFOSEQ
DPACOMPUTERSEQ	DPAMPROC VARIANTSEQ	OMPCIRLNSEQ
DPACPUSEQ	DPAMSOFTWARESEQ	

4. Run the SQL script on the database to change the sequence cache values.

Related reference:

“Optimized access to data” on page 105

The use of sequence cache and separate table spaces for large tables can help you improve your system performance. Also, you can set applicable page sizes and storage capacities to these table spaces and cache memories for optimal performance.

Optimizing performance in DB2

You can change configuration settings at the registry, database manager, and the database levels. The registry level controls all DB2 databases and DB2 applications. The database manager controls the main configuration settings for all databases. Each individual database instance can also have its own settings at the database level.

Before you begin

To change settings, you must log in as the instance owner, which is the owner that created the specific instance, for example, db2inst1. In a Windows environment, the user who installed the database is the instance owner. In a UNIX or Linux environment, the root user installs the database and creates the instance owner. After installation of the database, only the instance owner can change the database settings.

About this task

Optimal performance is subjective and dependent on your environment and the needs and expectations of your users. In general, optimal performance typically includes response times within a tested and acceptable limit and the optimal uses of system resources. Hardware use should be efficient, which means that you do not have unused processing cycles or overworked system resources.

Any values for database settings are offered as a starting point, which you can refine to meet the needs of your deployment. Values were determined through performance testing in several different environments, which include the AIX®, Windows, and Linux operating systems.

To determine the optimal performance settings for your environment, conduct performance tests before deployment.

Related tasks:

Manually configuring DB2 9.x

“Developing performance tests” on page 130

Performance tests provide baseline measurements of system performance, which you can use to identify the weaknesses in your deployment and to determine optimal configuration settings. When you design performance tests, you determine test objectives, develop use cases, develop a test strategy, and define your test environment.

Related information:



IBM DB2 Information Center (Opens in a new browser window or tab.)

High availability for databases

Setting environment variables and registry variables for optimal performance

Registry variables store configuration settings that can be applied at a global or instance level. An environment variable can set multiple registry variables based on pre-defined settings. When you set the **DB2_WORKLOAD** environment variable to **MAXIMO**, you efficiently change several registry variables to settings that are optimized for Maximo Asset Management.

About this task

The **MAXIMO** value for the **DB2_WORKLOAD** environment variable was first introduced in DB2 version 9.5, fix pack 5. If you are using DB2 version 9.5 do not have fix pack 5, you cannot use the **MAXIMO** value for the **DB2_WORKLOAD** environment variable.

Procedure

1. In the command prompt on the server on which DB2 is installed, run the following command:

```
db2set DB2_WORKLOAD=MAXIMO
```

The following registry values are automatically set:

Registry variable	Value
DB2_SKIPINSERTED	ON
DB2_INLIST_TO_NLJN	YES
DB2_MINIMIZE_LISTPREFETCH	YES
DB2_EVALUNCOMMITTED	YES
DB2_SKIPDELETED	ON

2. If you are using DB2 version 9.7, or later, run the following command:
db2set DB2_USE_ALTERNATE_PAGE_CLEANING=ON The **DB2_USE_ALTERNATE_PAGE_CLEANING** registry variable speeds up the process of creating or altering large table spaces.
3. If you are using the Windows or UNIX operating system, run the following command:
db2set DB2_FMP_COMM_HEAPSZ=65536 The **DB2_FMP_COMM_HEAPSZ** registry variable sets the size of the heap that is used for fenced routine invocations.
4. To apply the settings that you changed, stop and restart DB2.

Related reference:

“DB2 registry variables”

The configuration settings that are stored in registry variables can be applied at a global or instance level. The registry variables that affect performance involve the time required to process commands, the optimization of SQL queues, the behavior of row locking, and heap sizes.

DB2 registry variables

The configuration settings that are stored in registry variables can be applied at a global or instance level. The registry variables that affect performance involve the time required to process commands, the optimization of SQL queues, the behavior of row locking, and heap sizes.

The following values are provided for tuning the registry variables in DB2:

Table 26. DB2 values to use when you tune the registry values that affect performance

Registry variables	Starting value	Purpose
DB2_SKIPINSERTED	ON	Decreases the time required to process certain database commands because uncommitted inserted rows are skipped.
DB2_INLIST_TO_NLJN	YES	Optimizes the performance of SQL queries that use the IN predicate.
DB2_MINIMIZE_LISTPREFETCH	YES	Prevents the use of list prefetch when prefetching would not be a good access method, for example, when the catalog statistics are not available. List prefetch is a special table access method that retrieves information from the index, sorts by page number, and then prefetches the pages.
DB2_EVALUNCOMMITTED	YES	Minimizes row locking until the table or index access scan determines that the data record matches the query predicate.
DB2_SKIPDELETED	ON	Decreases the time required to process database commands because deleted rows are skipped.
DB2_USE_ALTERNATE_PAGE_CLEANING	ON	Speeds up the process of creating or altering large table spaces and database restore operations.
DB2_FMP_COMM_HEAPSZ	65536	Sets the size of the heap that is used for fenced routine invocations, such as stored procedures or user-defined function calls. The size is measured in 4-KB pages.

Tuning database manager settings

The database manager settings that affect performance set the sizes of memory spaces, control the behavior of fenced processes, and specify whether the instance is monitored.

Procedure

1. For all operating systems, in the command prompt on the server on which DB2 is installed, run the following commands:
db2 update dbm cfg using RQRIOBLK 65535
db2 update dbm cfg using HEALTH_MON OFF
db2 update dbm cfg using MON_HEAP_SZ AUTOMATIC
db2 update dbm cfg using KEEPFENCED NO
2. Set the amount of virtual memory that is allocated for each agent:

Option	Description
For the UNIX or Linux operating system	Run the following command: db2 update dbm cfg using AGENT_STACK_SZ 1024
For the Windows operating system	Run the following command: db2 update dbm cfg using AGENT_STACK_SZ 1000

Related reference:

“DB2 database manager settings”

The database manager settings are set at the instance level. The settings that affect performance involve memory size, how fenced processes behave, and whether the database instance is monitored.

DB2 database manager settings

The database manager settings are set at the instance level. The settings that affect performance involve memory size, how fenced processes behave, and whether the database instance is monitored.

The following values are provided for tuning the database manager settings in DB2:

Database manager settings	Value for performance tuning	Descriptions
RQRIOBLK	65535	Sets the size of the communication buffer between remote applications and their database agents on the database server.
HEALTH_MON	OFF	Specifies whether an instance, the associated databases, and database objects are monitored.

Database manager settings	Value for performance tuning	Descriptions
MON_HEAP_SZ	AUTOMATIC	Specifies the amount of the memory, in 4-KB pages, that is allocated to database system monitor data. The AUTOMATIC value means that the monitor heap size can increase as needed until the limit of the instance memory is reached.
KEEPFENCED	NO	Indicates whether a fenced mode process is kept after a fenced mode routine call is complete. If you have many fenced processes, you can set the KEEPFENCED setting to YES, but then you must monitor your memory usage.
AGENT_STACK_SZ	For UNIX or Linux operating systems: 1024 For Windows operating system: 1000	Determines the amount of virtual memory, measured in 4-KB pages, that is allocated for each agent.

Tuning database configuration settings

You can specify database configuration settings for each individual database in DB2. Several database configuration settings can affect performance, such as the log file sizes or asynchronous cleaner options.

Procedure

- For all operating systems, in the command prompt on the server on which DB2 is installed, run the following commands:

```
db2 update db cfg for dbname using CHNGPGS_THRESH 40
db2 update db cfg for dbname using DFT_QUERYOPT 5
db2 update db cfg for dbname using LOGBUFSZ 1024
db2 update db cfg for dbname using LOGFILSIZ 8096
db2 update db cfg for dbname using LOGPRIMARY 20
db2 update db cfg for dbname using LOGSECOND 100
db2 update db cfg for dbname using LOCKLIST AUTOMATIC
db2 update db cfg for dbname using LOCKTIMEOUT 300
db2 update db cfg for dbname using NUM_IOCLEANERS AUTOMATIC
db2 update db cfg for dbname using NUM_IOSERVERS AUTOMATIC
db2 update db cfg for dbname using SOFTMAX 1000
db2 update db cfg for dbname using PCKCACHESZ 524288
db2 update db cfg for dbname using STAT_HEAP_SZ 51200
```
- Set the maximum number of file handles that can be open per application.

Operating system	Commands
UNIX or Linux	db2 update db cfg for <i>dbname</i> using MAXFILOP 61440
Windows	db2 update db cfg for <i>dbname</i> using MAXFILOP 65535

- If you recently upgraded to DB2 version 9.7, ensure that the following database configuration settings are set to the following values:

Setting	Value
CUR_COMMIT	ON
AUTO_REVAL	DEFERRED
DEC_TO_CHAR_FMT	NEW
STMT_CONC	LITERALS

What to do next

When the **STMT_CONC** parameter is set to LITERALS, the statement concentrator is enabled. All queries are translated into queries with parameter markers, which are used by the query optimizer when selecting an access plan. To assist the query optimizer to select an efficient access plan, enable the REOPT(ONCE) bind option.

Related tasks:

“Enabling the REOPT(ONCE) bind option”

To help the query optimizer select an efficient access plan, specify the REOPT(ONCE) bind option when you run queries. When the REOPT(ONCE) bind option is used, the query optimizer selects the access plan the first time that the query is run. Each subsequent time that the query is run, the access plan is reused.

Related reference:

“DB2 database configuration settings” on page 113

Each database in DB2 has a separate set of configuration settings. The database configuration settings that effect performance define log file sizes, memory sizes, asynchronous cleaner options, locklist sizes, and the maximum number of file handlers.

Related information:



Performance Tip: REOPT(ONCE) when using DB2 Statement Concentrator (Opens in a new browser window or tab.)



DB2 9.7 Statement Concentrator with STMM (Opens in a new browser window or tab.)

Enabling the REOPT(ONCE) bind option

To help the query optimizer select an efficient access plan, specify the REOPT(ONCE) bind option when you run queries. When the REOPT(ONCE) bind option is used, the query optimizer selects the access plan the first time that the query is run. Each subsequent time that the query is run, the access plan is reused.

Procedure

1. On the DB2 server, log in as the database instance owner.
2. In the bnd directory for DB2, for example C:\DB2\SQLLIB\bnd, run the following command:
db2 bind db2clipk.bnd collection NULLIDR1
3. Log in to Maximo Asset Management as an administrator with the authority to perform system configuration tasks.
4. In the System Properties application, open the details for the **mxe.db.DB2jdbcCollection** system property.
5. In the Global Properties Details section, in the **Global Value** field, specify NULLIDR1.

6. From the **Select Action** menu, select **Live Refresh**, and then click **OK**. The new property value, NULLIDR1, is shown in the **Global Value** field and the **Current Value** field.

Related information:

 Using the REOPT bind option with input variables in complex queries (Opens in a new browser window or tab.)

 Performance Tip: REOPT(ONCE) when using DB2 Statement Concentrator (Opens in a new browser window or tab.)

DB2 database configuration settings

Each database in DB2 has a separate set of configuration settings. The database configuration settings that effect performance define log file sizes, memory sizes, asynchronous cleaner options, locklist sizes, and the maximum number of file handlers.

The following values are provided for tuning the database configuration settings in DB2:

Setting	Starting value	Purpose
CHNGPGS_THRESH	40	Specifies the percentage of changed pages after which the asynchronous page cleaners are started.
DFT_QUERYOPT	5	Provides significant query optimization with heuristics to limit the effort used to select an access plan.
LOGBUFSZ	1024	Specifies the amount, in 4-KB pages, of the database heap to use as a buffer for log records before the records are written to disk.
LOGFILSIZ	8096	Defines the size of the log files, in 4-KB pages.
LOGPRIMARY	20	Sets the number of primary log files, which establish a fixed mount of storage that is allocated to the recovery log files.
LOGSECONDARY	100	Sets the number of secondary log files that are created and used for recovery log files when needed.

Setting	Starting value	Purpose
LOCKLIST	AUTOMATIC	Specifies the amount of storage that is allocated to the lock list, which contains the locks held by all applications concurrently connected to the database. The AUTOMATIC value means that when the workload requirements change, the memory tuner is able to dynamically size the storage for the lock list.
LOCKTIMEOUT	300	Specifies the number of seconds that the database waits before locking.
NUM_IOCLEANERS	AUTOMATIC	Sets the number of asynchronous page cleaners for the database. The AUTOMATIC value means that the number of page cleaners is based on the number of processors configured on the current server and the number of local logical database partitions in a partitioned database environment.
NUM_IOSERVERS	AUTOMATIC	Sets the number of I/O servers, which runs the prefetching operation and utilities. The AUTOMATIC value means that the number of prefetchers is calculated at database activation time.
SOFTMAX	1000	Determines the frequency of soft checkpoints and the recovery range. This setting is measured in the percentage of the size of one primary log file.
STMHEAP	20000	Sets the size of the statement heap, which is used as the workspace for the SQL compiler.
MAXFILOP	For UNIX or Linux operating systems: 61440 For Windows operating system: 65535	Sets the maximum number of file handles that can be open per application.
CUR_COMMIT	ON	Controls the behavior of cursor stability scans.

Setting	Starting value	Purpose
AUTO_REVAL	DEFERRED	Controls the revalidation and invalidation semantics. The DEFERRED setting means that all dependent objects are revalidated at the time of next access.
DEC_TO_CHAR_FMT	NEW	Controls the result of the CHAR scalar function and the CAST specification for converting decimal to character values. The NEW setting means that leading zeros and trailing decimal characters are not included in the result of the CHAR function.
STMT_CONC	LITERALS	Enables and sets the default statement concentrator behavior. The LITERALS setting means that SQL statements that are identical, except for the values of literals in the statements, can share package cache entries. After you enable the statement concentrator, do not use the VARGRAPHICS function. To assist the query optimizer to select an efficient access plan, enable the REOPT(ONCE) bind option.
PCKCACHESZ	524288	Specifies the amount of database shared memory, measured in 4-KB pages, that caches sections for static and dynamic SQL and XQuery statements on a database.
STAT_HEAP_SZ	51200	Sets the maximum size of the heap, measured in 4-KB pages, that is used to collect statistics when the RUNSTATS command is run.

Related tasks:

“Enabling the REOPT(ONCE) bind option” on page 112

To help the query optimizer select an efficient access plan, specify the REOPT(ONCE) bind option when you run queries. When the REOPT(ONCE) bind option is used, the query optimizer selects the access plan the first time that the query is run. Each subsequent time that the query is run, the access plan is reused.

Related information:

 Performance Tip: REOPT(ONCE) when using DB2 Statement Concentrator (Opens in a new browser window or tab.)

 DB2 9.7 Statement Concentrator with STMM (Opens in a new browser window or tab.)

Reorganization of tables and indexes in DB2

When many updates are made on a table, the space can become fragmented. You can reorganize data to reclaim space in a table and to improve system performance. Use the IBM DB2 **REORGCHK** and **REORG** commands to optimize table spaces and indexes. You cannot reorganize metadata tables.

The following conditions can indicate that the database tables require reorganization:

- The tables experience a high volume of insert, update, and delete activity.
- You observe a decrease in the performance of queries.
- You run the **RUNSTATS** command to refresh the table statistics but performance does not improve.
- You run the **REORGCHK** command, and the results indicate that table reorganization is required.

The following conditions can indicate that the indexes require reorganization:

- The leaf pages, which contain the pairs of keys and record identifiers that point to the actual data, are fragmented. When leaf pages are fragmented, performance is affected because more leaf pages must be read to fetch table pages.
- The physical index page does not match the sequence of keys on the pages, which increases the time required to process queries.
- The index has too many levels.

If you determine that a reorganization can improve the performance of the table or index, you can complete the reorganization online or offline. By default, the **REORG** command completes offline reorganizations.

Database maintenance can reduce the time needed for reorganization. Your database maintenance routine should include archiving and trimming tables on a regular basis. Your regular database maintenance should also include running the **REORGCHK** command to identify tables and indexes that might require reorganization.

Related concepts:

“Database indexing” on page 104

Indexing a database requires good understanding of the data, the user functions, and how the database is indexed. Indexes use key parts of data from a table in a binary structure to improve searching capability. Each record of data in the table must be associated with data in the index.

Related information:



Table reorganization (Opens in a new browser window or tab.)



Index reorganization (Opens in a new browser window or tab.)



REORGCHK command (Opens in a new browser window or tab.)



REORG TABLES command (Opens in a new browser window or tab.)

DBMS_STATS package

The DBMS_STATS package in Oracle optimizes statistics on your database. The system benefits from cost-based optimization because it builds many queries dynamically, depending on user input.

Optimizing performance in Oracle Database

Initialization parameters set values that can affect system performance, such as the optimizer features. Initialization parameters are stored in an initialization parameter file. You can change initialization parameters with **ALTER SYSTEM** commands.

Before you begin

To set parameters, you must log in as a user that has SYSDBA privileges.

Procedure

1. Set the **CURSOR_SHARING** parameter to SIMILAR or FORCE so that literal values are converted to bind variables.
2. If the database character set is a double-byte or Unicode character, set the **NLS_LENGTH_SEMANTICS** parameter to CHAR.
3. Set the **WORKAREA_SIZE_POLICY** parameter to AUTO to automatically size the work areas.
4. Ensure that the **OPTIMIZER_FEATURES_ENABLE** parameter is set to your current version of Oracle Database.
5. Set the **PROCESSES** parameter to the maximum number of users and background processes that can access the database concurrently. For example, if you expect 50 concurrent users, set the value to 70 to include background processes.
6. Set the **OPEN_CURSORS** parameter to the maximum number of open cursors that a session can have at one time. Open cursors handle private SQL areas. The number of open cursors that you require depends on your deployment. Set the value high enough to prevent Maximo Asset Management from running out of cursors.
7. Set the **SESSIONS** parameter to specify the maximum number of sessions that can be created. Set the **SESSIONS** parameter to a value that is based on the maximum number of users and background processes plus an allowance of

10% for recursive processes. For example, if you expect 50 concurrent users, set the value to 77, which can accommodate 20 background processes with seven sessions for recursive processes.

8. Set the System Global Area (SGA) and Program Global Area (PGA) management parameters to a memory size that is based on the database size, number of concurrent users, and workload.

Oracle Database version	Parameters to set
For Oracle Database 10g	SGA_TARGET SGA_MAX_SIZE
For Oracle Database 11g	SGA_TARGET SGA_MAX_SIZE MEMORY_TARGET MEMORY_MAX_TARGET

9. Set the **TRANSACTIONS** parameter to specify the maximum number of concurrent transactions. A larger value for this parameter means that the size of the SGA is also larger.

10. To apply your changes, stop, and restart Oracle Database.

Related tasks:

“Developing performance tests” on page 130

Performance tests provide baseline measurements of system performance, which you can use to identify the weaknesses in your deployment and to determine optimal configuration settings. When you design performance tests, you determine test objectives, develop use cases, develop a test strategy, and define your test environment.

Related reference:

“Oracle Database initialization parameters”

Initialization parameters are stored in an initialization parameter file and can be applied to all database instances on a server. Parameters that affect system performance involve cursor sharing, the policy that determines work area size, the number of concurrent processes, and memory area sizes.

Related information:

 [Oracle Database Performance Tuning Guide 10g \(Opens in a new browser window or tab.\)](#)

 [Oracle Database Performance Tuning Guide 11g \(Opens in a new browser window or tab.\)](#)

[High availability for databases](#)

Oracle Database initialization parameters

Initialization parameters are stored in an initialization parameter file and can be applied to all database instances on a server. Parameters that affect system performance involve cursor sharing, the policy that determines work area size, the number of concurrent processes, and memory area sizes.

Setting	Starting value	Purpose
CURSOR_SHARING	SIMILAR or FORCE	Determines which SQL statements can share cursors. With the SIMILAR value or the FORCE value, SQL statements share cursors when differences between the statements do not affect the meaning or optimization of the statement.
NLS_LENGTH_SEMANTICS	CHAR	Creates char and varchar2 columns that use character length semantics.
WORKAREA_SIZE_POLICY	AUTO	Automatically sizes the work areas that are used by memory-intensive operators.
OPTIMIZER_FEATURES_ENABLE	The current release number for the version of Oracle Database that you are using, such as 11.1.0.7	Enables a series of features that optimize performance based on an Oracle release number
SGA_TARGET	A numerical value for the memory size that is based on the database size, number of concurrent users, and workload.	In Oracle Database 10g, and Oracle Database 11g, sets the total size of all System Global Area (SGA) memory. When this parameter is set, the buffer cache, Java pool, large pool, and shared pool settings are all sized automatically
SGA_MAX_SIZE	A numerical value for the memory size that is based on the database size, number of concurrent users, and workload.	In Oracle Database 10g, and Oracle Database 11g, specifies the maximum size of all SGA memory.
MEMORY_TARGET	A numerical value for the memory size that is based on the database size, number of concurrent users, and workload.	In Oracle Database 11g, sets the system global area (SGA) or program global area (PGA) memory sizes. The SGA and PGA are dynamically sized as needed, based on this setting.
MEMORY_MAX_TARGET	A numerical value for the memory size that is based on the database size, number of concurrent users, and workload.	In Oracle Database 11g, sets the maximum value that can be used in the MEMORY_TARGET parameter.
PROCESSES	A numerical value that is based on the number of concurrent users that you expect to connect to the database. Add more processes for background processes. For example, if you expect 50 concurrent users, then you might set the PROCESSES parameter to 70.	Sets the maximum number of processes, both user and background processes, that can concurrently connect to the database.

Setting	Starting value	Purpose
OPEN_CURSORS	A numerical value that sets the number of open cursors that are available.	Sets the number of open cursors, which handle private SQL areas.
SESSIONS	A numerical value that is based on the number of concurrent users that you expect to connect to the database. Add more sessions for background processes and a 10% allowance for recursive processes. For example, if you expect 50 concurrent users, set the value to 77, which can accommodate 20 background processes with 10%, or 7 sessions, for recursive processes.	Sets the maximum number of sessions that can be created.
TRANSACTIONS	A numerical value that specifies the maximum number of concurrent transactions.	Sets the maximum number of concurrent transactions. When this value is set to a higher number, the SGA is larger.

Related information:

 [Oracle Database Performance Tuning Guide 10g](#) (Opens in a new browser window or tab.)

 [Oracle Database Performance Tuning Guide 11g](#) (Opens in a new browser window or tab.)

IBM WebSphere Application Server performance tuning

You can define initial and maximum heap sizes and thread pool settings. You can tune the Java virtual machine (JVM) for optimal performance by setting up the JVM parameters.

You must tune the application server at the JVM level because the memory usage and the garbage collection options affect each JVM separately. These settings are used for tuning the JVM performance in Tivoli test environments, and so your environment might require different settings. You can use these JVM settings as a guideline or as a starting point, and then configure these JVM settings to your environment requirements.

Related concepts:

Performance improvements with clusters

A cluster groups similar functions on two or more Java virtual machines (JVMs) to process a single function, such as scheduled cron tasks. Clusters connect to the same database but operate independently. For example, if the cron task cluster fails, users can still connect to the user interface cluster.

Related tasks:

“Developing performance tests” on page 130

Performance tests provide baseline measurements of system performance, which you can use to identify the weaknesses in your deployment and to determine optimal configuration settings. When you design performance tests, you determine test objectives, develop use cases, develop a test strategy, and define your test environment.

Related information:

 [Monitoring performance with Tivoli Performance Viewer \(TPV\)](#) (Opens in a new browser window or tab.)

[Configuring application servers for high availability](#)

 [Tuning the IBM virtual machine for Java](#) (Opens in a new browser window or tab.)

[High availability for application servers](#)

 [Case study: Tuning WebSphere Application Server V7 and V8 for performance](#) (Opens in a new browser window or tab.)

 [Solving memory problems in WebSphere applications](#) (Opens in a new browser window or tab.)

 [Tuning WebLogic Server](#) (Opens in a new browser window or tab.)

Thread pool sizes

If you tune thread pool sizes, the server components can reuse threads. The reuse of threads eliminates the need to create new threads at run time to handle each new request.

To tune the thread pool size settings, you can use the **Thread pools** option in the WebSphere Application Server administrative console.

The default thread pool is used when requests arrive for message-driven beans or if a transport chain was not defined for a thread pool. The WebContainer thread pool is used when requests come over HTTP.

You can tune the following settings to improve performance in your system:

Minimum size

The minimum number of threads to maintain in the thread pool.

Default thread pool: 20

WebContainer thread pool: 120

Maximum size

The maximum number of threads to maintain in the thread pool.

Default thread pool: 50

WebContainer thread pool: 120

Thread inactivity timeout

The amount of inactivity (in milliseconds) that can elapse before a thread is reclaimed.

Default thread pool: 30000

WebContainer thread pool: 60000

Allow thread allocation beyond maximum thread size

If enabled, the number of threads can increase beyond the maximum size configured for the thread pool.

Value: Enabled

Related information:

Thread pool settings (Opens in a new browser window or tab.)

Heap size values

Java virtual machine (JVM) heap size parameters directly influence garbage collection behavior. If you increase the heap size value, your system can process more objects before the heap size triggers a garbage collection.

The tuning of JVM heap sizes often involves a balance between the garbage collections and the time needed to perform the garbage collection. A larger heap size also means that a larger amount of time is needed to find and process objects that need to be collected.

When you plan for system memory consumption, include additional processor memory for the JVM to use outside of the heap size and random access memory (RAM) for the operating system. Include an additional 30% or 40% of memory to account for this additional processor usage.

The value of the JVM heap size is directly related to the amount of physical memory in the system. Set the initial and maximum heap sizes to 4096 to start tuning, because 64-bit operating systems have an address space limit of 4 GB, regardless of the amount of physical memory in the system.

Never set the JVM heap size larger than the physical memory in the system.

Related tasks:

“Determining optimal heap sizes in WebSphere Application Server”

The optimal heap size for your deployment ensures that memory is not wasted or constrained. To determine the optimal heap size, you enable verbose garbage collection, and then use a support tool to analyze the results and determine the optimal size.

Related information:



Handling out-of-memory situations (Opens in a new browser window or tab.)



Activating the heap monitor (Opens in a new browser window or tab.)



Memory and Address Space Limits (Opens in a new browser window or tab.)

Determining optimal heap sizes in WebSphere Application Server

The optimal heap size for your deployment ensures that memory is not wasted or constrained. To determine the optimal heap size, you enable verbose garbage collection, and then use a support tool to analyze the results and determine the optimal size.

Procedure

1. In the navigation pane of the WebSphere Application Server administrative console, select **Servers > Application servers > *server_name* > Process definition > Java Virtual Machine**.
2. Select the **Verbose Garbage Collection** check box. When verbose garbage collection is enabled, the Java virtual machine (JVM) records information about each garbage collection in a log file. For example, in the log file, you can see the amount of free bytes and used bytes in the heap, the interval between garbage collections, and the pause time. Verbose garbage collection has minimal effect on system performance.
3. For Oracle Solaris and HP-UX, add the following parameters to the generic JVM arguments:
 - XX:+PrintGCDetails
 - XX:+PrintGCTimeStamps
 - XX:+PrintHeapAtGC
4. Apply and save the changes.
5. Restart WebSphere Application Server.
6. To create log entries to analyze, allow the system to process a typical user load for a specified time. The time might be a few hours or a few days, depending on the user load.
7. To determine the optimal heap size, analyze the log file with a garbage collection analyzer. You can use the IBM Monitoring and Diagnostic Tools for Java - Garbage Collection and Memory Visualizer plug-in, which is available for IBM Support Assistant.
 - For AIX , Microsoft Windows, or Linux, analyze the `native_stderr.log` file.
 - For Oracle Solaris or HP-UX, analyze the `native_stdout.log` file.
8. Specify new initial and maximum heap sizes and save your changes.
9. Disable verbose garbage collection.
10. Delete the `.log` file. The `.log` file can grow large when verbose garbage collection is enabled.

Related concepts:

“Heap size values” on page 122

Java virtual machine (JVM) heap size parameters directly influence garbage collection behavior. If you increase the heap size value, your system can process more objects before the heap size triggers a garbage collection.

Related information:



IBM Monitoring and Diagnostic Tools for Java - Garbage Collection and Memory Visualizer (Opens in a new browser window or tab.)

JVM commands to optimize performance

The generic Java virtual machine (JVM) arguments are optional command-line arguments that are passed to the JVM when WebSphere Application Server starts. The generic JVM arguments can set the timeout value for the server-side Java remote method invocation (RMI), disable explicit garbage collection, set the garbage collection policy, and specify the nursery size.

When you configure the JVM, you can specify the optional command-line arguments in the **Generic JVM arguments** field. To use more than one argument, enter a space between each argument.

The following generic JVM arguments can improve system performance:

-sun.rmi.dgc.ackTimeout=10000

Sets the time in milliseconds that the server-side Java RMI runtime strongly refers to a remote object. Because RMI allocates a large quantity of short-lived remote objects, a value for the `sun.rmi.dgc.ackTimeout` argument that is too high can prevent the garbage collection from operating efficiently, which can cause out-of-memory problems. The `-sun.rmi.dgc.ackTimeout=10000` argument sets the value to 10000 (10 seconds), which can prevent out-of-memory problems.

-Xdisableexplicitgc

Disables explicit garbage collection, which prevents `System.gc()` calls from starting the garbage collection process.

-Xmn1024m

Sets the size of the nursery to 25% of the maximum heap size. The nursery is the area in the heap where objects are created. If you analyze the garbage collection and then adjust the heap sizes, adjust the nursery size to reflect your changes.

-Xgcpolicy:gencon

Sets the garbage collection policy to `gencon` garbage collection, which places objects in separate areas of the heap based on their lifetime. After objects are created in the nursery and then survive a number of garbage collections, the objects are moved to a tenured area. When objects are separated in this way, garbage collection can run more frequently in the nursery without affecting the rest of the heap, which keeps pauses to a minimum. Because Maximo Asset Management creates many short-lived objects, set the garbage collection policy to `gencon`.

Related information:

-  [Tuning the IBM virtual machine for Java \(Opens in a new browser window or tab.\)](#)
-  [Solving memory problems in WebSphere applications \(Opens in a new browser window or tab.\)](#)
-  [Java virtual machine settings \(Opens in a new browser window or tab.\)](#)

HTTP server performance tuning

You can tune the IBM HTTP Server to improve Maximo Asset Management response times. You can customize your parameter settings to improve performance.

You can use the IBM HTTP Server parameter settings to optimize tuning. These settings are used for tuning performance in a test environment, and so your environment might require different settings. You can use these settings as a guideline or as a starting point, and then monitor and tune the settings to your environment.

The versions of IBM HTTP Server and of WebSphere Application Server must be the same.

The following parameter settings were optimized and tested for the Microsoft Windows environment. Although these same settings were also tested and optimized in an AIX environment, they are not optimized for UNIX or Linux.

IBM HTTP Server working on a Windows environment has a parent process and a single, multithreaded child process which creates simultaneous connections. Set the following parameters so that your environment can handle the simultaneous connections more efficiently:

TimeOut

The amount of time that the server waits for certain events before a request is rejected.

Value: 900

KeepAliveTimeOut

The amount of time that the server waits for subsequent requests on a persistent connection. Set this value to 60 for those environments that have high network latency. For example, network latency can be an issue when users are in areas that are geographically different from the location of the servers. That means the further the users are located from the servers, the higher the network latency is.

High network bandwidth: 10

Low network bandwidth: 60

MaxKeepAliveRequests

The number of requests allowed on a persistent connection. If you limit the number of requests, the server must finish the connection and create another connection when the limit is reached. The need to finish and create connections continuously can affect server performance. Set this value to zero to allow unlimited requests on a persistent connection.

Value: 0

MaxRequestsPerChild

The limit on the number of requests that an individual child process handles during its lifecycle. After this limit is reached, the child process ends. When you set this value to 0, the child process never ends and can handle unlimited requests. If you set this value to a number greater than 0, extra processing is required to terminate and create child processes. On Linux and UNIX systems, a value other than 0 can create a high number of child processes, which can result in excessive swap space usage. To minimize potential issues, set this value to 0.

Value: 0

ThreadLimit

The upper limit on the configurable number of threads per child process.

Value for Windows environments: 2400

ThreadsPerChild

The number of threads that each child process can create. In a Windows environment running on 64 bits, each instance is limited to 2500 threads per child approximately. For the 32-bit environments, the total number of threads per child is closer to 5000. These numbers are not exact limits because the actual limits are the sum of the startup memory that is used by each thread plus the maximum runtime memory usage per thread. That sum varies based on configuration and workload. If you raise the number of the **ThreadsPerChild** parameter, you risk having child processes that do not work when the runtime memory raises the address space over 2 or 3 GB.

Value for Windows environments: 2400

You can set the **ThreadsPerChild** and **ThreadLimit** parameters to the same value.

Related tasks:

“Developing performance tests” on page 130

Performance tests provide baseline measurements of system performance, which you can use to identify the weaknesses in your deployment and to determine optimal configuration settings. When you design performance tests, you determine test objectives, develop use cases, develop a test strategy, and define your test environment.

Related information:

IBM HTTP Server performance tuning (Opens in a new browser window or tab.)

High availability for HTTP servers

IBM HTTP Server compression and load balancing

HTTP compression improves the usage of available bandwidth and provides faster transmission speeds. HTTP compression is built into web servers and web browsers. You can configure settings for HTTP compression and load balancing to achieve optimal performance.

HTTP compression affects the data on all servers in a cluster. Each compression-compliant browser receives the compressed data in the format that the browser supports. If the browser does not support compression, then data downloads in uncompressed format.

Data is compressed by using a compression module such as the `mod_deflate` module from Apache. The compression method is dictated by the software installed on the server.

In IBM HTTP Server, use the Apache `mod_deflate` module and set **DeflateCompressionLevel** to 3 or 6 to improve response time in environments that have low bandwidth and high latency.

Load balancing

Hardware load balancers that bypass IBM HTTP Server cannot employ the data compression method. If bypassing IBM HTTP Server, you must set up the load balancing to distribute the task load across multiple instances of an application. The user load comes from users who are logged in to the system and use the interface to complete tasks. Non-user load comes from items such as scheduled jobs (cron tasks) and transactions that come from Maximo Integration Framework.

You can distribute user load and non-user load across different application servers or clusters by using the IBM HTTP Server for WebSphere Application Server plug-in. The plug-in acts as an agent that uses the HTTP protocol to redirect HTTP requests from the web server to the application server.

To improve performance by using this plug-in, you modify the load balancing option. The plug-in uses this option to send requests to the application servers that are associated with the web server. Both load balance options, RoundRobin and Random, provide an even distribution of work across cluster members. However, you must test the options to determine which works better for your deployment.

To choose the best load balance option, you can configure Maximo Asset Management to log the number of users per JVM. The data can help you determine which option provides the best load balancing.

Related information:

Utilities for logging and testing

To help you manage system performance, there are utilities that you can use for testing and debugging purposes. When you are satisfied with your deployment, you can disable these logging utilities.

Topology 1: Load-balanced configuration model

Configuring HTTP servers for high availability

 [Handling enough simultaneous connections with IBM HTTP Server 2.0 and above on Linux and Unix systems \(Opens in a new browser window or tab.\)](#)

Optimized settings for operating system configuration

Maximo Asset Management is supported by several operating systems, such as AIX, Red Hat Enterprise Linux, and Windows. You can configure the system settings at the operating level on application and database servers for optimal performance.

You must check with your network administration group to make sure that the suggested setup values are compatible with the installed network. Setup values must also conform with the standard settings configured by your network administration team.

Related tasks:

“Developing performance tests” on page 130

Performance tests provide baseline measurements of system performance, which you can use to identify the weaknesses in your deployment and to determine optimal configuration settings. When you design performance tests, you determine test objectives, develop use cases, develop a test strategy, and define your test environment.

Related information:



Performance Tip: Understanding Operating System Tuning Parameters (Opens in a new browser window or tab.)

Performance-related settings on AIX

The AIX operating system can be configured to improve Maximo Asset Management performance. These settings include improved setup of network, resources, processing, and virtual memory.

You must apply the network, process, and virtual memory settings to the root user of the WebSphere Application Server instance owner for application servers and to the root user of the database instance owner for the database server. When you apply resource (ulimit) settings, you must apply the settings to the *userid* of the instance owner in both servers, but the root user is still required to make the changes.

Network settings

From the command line of the operating system, enter the following settings to optimize the network services:

```
/usr/sbin/no -r -o sb_max=6192000
/usr/sbin/no -r -o tcp_sendspace=4096000
/usr/sbin/no -r -o tcp_recvspace=4096000
/usr/sbin/no -r -o udp_sendspace=65536
/usr/sbin/no -r -o udp_recvspace=655360
/usr/sbin/no -r -o rfc1323=1
/usr/sbin/no -r -o ipqmaxlen=250
/usr/sbin/no -r -o clean_partial_conns=1
/usr/sbin/no -r -o tcp_keepidle=600
/usr/sbin/no -r -o tcp_keepintvl=10
/usr/sbin/no -r -o tcp_keepinit=40
/usr/sbin/no -r -o tcp_timewait=1
/usr/sbin/no -r -o tcp_finwait2=60
/usr/sbin/no -r -o tcp_ephemeral_low=1024
```

Note that the values for the **sb_max**, **tcp_sendspace**, and **tcp_recvspace** parameters depend on the device type and speed. You must check the values of your TCP streaming workload to find out the best tuning value for your network.

These network settings apply to the whole system, except for the ulimit settings on AIX.

Resource (ulimit) settings

The ulimit settings are changed for specific users, such as the DB2 instance owner, the DB2 fence owner, the Maximo Asset Management user, and the WebSphere Application Server admin user. Run the following command line to set up the

resource (ulimit) settings of your operating system resources:
chuser fsize=-1 fsize_hard=-1 data=-1 data_hard=-1 stack=4194304
stack_hard=4194304 nofiles=-1 nofiles_hard=-1 <user_name>

Process settings

From the command line of your operating system, enter the `chdev -l sys0 -a maxuproc='4096'` process setting.

Virtual memory settings

To set the tunable parameters for the Virtual Memory Manager, enter the following settings on your command line:

```
vmo -p -o lru_file_repage = 0  
vmo -p -o maxclient% = 90  
vmo -p -o maxperm%=90  
vmo -p -o minperm%=5
```

Related information:

-  [TCP streaming workload tuning \(Opens in a new browser window or tab.\)](#)
-  [vmo Command \(Opens in a new browser window or tab.\)](#)
-  [no Command \(Opens in a new browser window or tab.\)](#)
-  [chuser Command \(Opens in a new browser window or tab.\)](#)
-  [Performance Tip: Understanding Operating System Tuning Parameters \(Opens in a new browser window or tab.\)](#)

Performance-related network parameters for Windows and Red Hat Enterprise Linux

You can configure the Windows or the Red Hat Enterprise Linux operating system to improve the application performance. You must set the network parameters to match the Maximo Asset Management requirements.

Windows parameters

You can set up the following network parameters, located under the Windows registry key, to improve system performance.

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters:  
  TcpTimedWaitDelay dword:0000001e (30)  
  StrictTimeWaitSeqCheck dword:00000001 (1)  
  MaxFreeTcbs dword:00011940 (72000)  
  MaxHashTableSize dword:0000ffff (65535)  
  TcpWindowSize dword:0000ffff (65535)  
  EnabledDynamicBacklog dword:00000001 (1)  
  MinimumDynamicBacklog dword:00000032 (20)  
  MaximumDynamicBacklog dword:000003eb (1000)  
  DynamicBacklogGrowthDelta dword:0000000a (10)  
  Interfaces\TcpAckFrequency dword:00000001 (1)  
  MaxUserPort dword:0000ffff (65535)
```

These settings apply to Microsoft Windows Server 2003. For Microsoft Windows Server 2008, the default dynamic port range is changed. The new default start port is 49152 and the default end port is 65535. Therefore, 16,384 ports are available by default (not 5,000).

To view the dynamic port range, start the command prompt and use the **netsh** command:

```
netsh int ipv4 show dynamicport tcp
```

To change the dynamic port range for the maximum number of ports allowed, run the following command:

```
netsh int ipv4 set dynamicport tcp start=1025 num=64510
```

The minimum start port is 1025 and the maximum end port cannot exceed 65535.

Red Hat Enterprise Linux parameters

For networking, enter the following command to improve your application performance:

```
sysctl -w net.ipv4.ip_local_port_range="1024 65535"
```

Related information:

 [Performance Tip: Understanding Operating System Tuning Parameters \(Opens in a new browser window or tab.\)](#)

Developing performance tests

Performance tests provide baseline measurements of system performance, which you can use to identify the weaknesses in your deployment and to determine optimal configuration settings. When you design performance tests, you determine test objectives, develop use cases, develop a test strategy, and define your test environment.

Related information:

 [Performance test best practices using Rational Performance Tester white paper \(Opens in a new browser window or tab.\)](#)

Determining test objectives

Performance tests can validate sizing estimates, ensure that your deployment meets your business or organization requirements, and address performance issues. You must identify the key business questions for your organization and develop tests that answer those questions.

Procedure

1. Identify the key questions that you want to answer in the performance tests.

For example, you might consider the following questions:

- Does the planned system architecture meet our business or organization requirements?
- Does the planned deployment provide a satisfactory response time for the expected number of concurrent users to perform a specified number of transactions over specific timeframe?
- Will a particular component create system performance issues for an expected transaction volume and number of concurrent users?
- Is the hardware in the deployment enough to provide acceptable performance, given the expected number of transactions and users?

2. Prioritize the risks, rewards, and costs in your deployment.

For example, you might decide that supporting large user loads that perform many concurrent transactions is the most important factor for your deployment.

Developing use cases

When you design performance tests, your goal is to develop few use cases that adequately test the functions that are most frequently used and most important.

Procedure

1. Determine how and when users log in and log off in the production environment. For example, do users stay logged in all day? Or do users log in, complete some transactions, and then log off?
2. Identify workloads that reflect the workload distribution and the workload rate. Distribution is the percentage of users who are completing a specific task, such as creating service requests. The rate is the transaction rate for a specific task, such as the number of service requests. Develop use cases that reflect both the number of transactions and the frequency of transactions. For example, 20% of users are creating 10 service requests in an hour.
3. Determine how user load changes throughout the course of the day. If you have real-time usage history, use the data for requests processed per second during peak hours to create your use case. Then, you can compare your real-time data with the server load created during the performance test to see how your deployment is functioning.
4. Based on user behavior, workloads, and user loads, create use cases that focus on your test objectives.

Developing test strategies

The type of performance test that you use depends on your test objectives. Focus your test strategy on identifying the type of performance test that can provide data to measure the objectives.

Procedure

1. Run tests to record baseline measurements for your deployment. A baseline measurement test is a single user test that is run over several iterations to identify whether an application is performing well. You can use baseline measurement data for comparison purposes in future tests.
2. Identify which performance tests can provide data that is related to your test objectives.

The type of tests that you might use include the following performance tests:

Benchmark under load test

A type of baseline measurement that increases the load from a single user to a percentage of your expected system load, such as 25%. This test indicates whether a system performance issue must be fixed before further testing.

Performance load tests

A multipart test that measures the response times of transactions during certain percentages of user loads. A typical performance load test consists of five cycles, each of which increases the percentage of the expected user load. For example, a performance load test can have five cycles at the following user load intervals: 50%, 75%, 100%, 125%, and

150%. Performance load tests establish the performance curve and identify whether the deployment can support service level agreements under user load.

Endurance tests

Tests that are run over a period of hours or days to identify memory leaks, performance issues over time, and overall system stability. During endurance tests, you monitor key performance indicators, such as transaction response times and the stability of memory usage.

Sizing and capacity tests

A series of tests that identifies the required size and total capacity at each tier of the deployment. Tiers can include the individual Java virtual machines (JVM), the total number of required JVMs, and the processor. You can use the results of sizing and capacity tests to determine the required resources for your deployment.

Batch tests

A test of components that does not require user interaction, such as cron tasks or integrations with external systems.

3. Identify the minimum number of tests that can produce the data required to answer the test objectives. Consider the time and resources that are available. A comprehensive and focused test strategy can produce results in a cost effective manner.

Defining test environments

Your performance environment is composed of many different components, such as operating system, middleware, and deployment topology. Plan the specifics of each component level to create a test environment.

About this task

A dedicated test environment, including both server and network components, produces the most reliable results. For example, bandwidth test numbers can be skewed by general traffic on the local area network. Select your monitoring tools to minimize effects on system performance.

Procedure

1. Ensure that your test environment meets the following requirements during benchmark tests:
 - An overall architecture that matches the production environment, such as the same operating system and middleware platforms, similar hardware proportions, and the same number of Java virtual machines (JVM).
 - The same versions of all deployed software.
 - Comparable, sufficient data in the databases. For example, test results can vary significantly if a query runs on a test database of 1000 records when your production database contains 50,000 records.
 - Identical server configurations. In the course of testing, you might find it necessary to modify the test server configuration and rebuild and deploy new enterprise archive (EAR) files. Keep copies of the previous EAR files and document any changes that you make.
2. Record the following configuration details for the servers in both the production and test environments:
 - Number of processors
 - Capacity or clock speed of processors

- RAM capacity
- Disk capacity
- Free space available on disks
- Network interface card (NIC) capacity
- Network bandwidth

What to do next

Write your test cases, then run your tests. After the tests are completed, analyze the test results.

Scenario: Developing performance tests to measure processor utilization

Company XYZ plans to deploy Maximo Asset Management with extensive customization. To ensure a successful deployment, Company XYZ develops and runs performance tests.

Background

Company XYZ plans to use Maximo Asset Management for asset management, purchasing, work order tracking. Because of specific business processes, Company XYZ has a customized deployment that uses automated workflows. The users in Company XYZ use the following applications:

- Assets
- Purchase Requisitions
- Purchase Orders
- Work Order Tracking

Step One: Determine the objectives to measure

The deployment team at Company XYZ considers the key business question and prioritizes the risks, rewards, and costs in the deployment. Based on research, the deployment team determines that their users do not like when transactions in web applications take too long to respond. Interviews with focus groups identify that users become frustrated when a transaction takes more than 2 seconds to respond.

The server management team at Company XYZ determines that if processor utilization remains below 80% for a target user load on the system, then the processor can provide adequate resources for applications to function at the wanted level. This value also can handle occasional spikes in processing, such as during month-end processing, without an effect on response times. Based on the size of the company, the server management team identifies 950 users as the target concurrent load.

Step Two: Develop use cases

The deployment team considers how users behave throughout the day. The team identifies that users generally log in after they arrive in the morning. Users typically complete a set of work activities and then log out. The deployment team estimates that each user completes a use case approximately 20 times per hour.

To approximate the login and logout behavior, the deployment team plans to create use cases in which automated test users log in, run six iterations of a use case, and

then log out. The automated test users then log in again and repeat the cycle. A pause of 5 to 10 seconds is incorporated into steps in the scripts to represent actual user processing rates.

The deployment team identifies the uses cases that are required to test the deployment. The team also assigns weight factors to each use case. The weight factor represents the number of automated users to run each use case.

Table 27. Uses cases identified for testing in Company XYZ

Use case identifier	Description	Weight factor
AS01	Search assets and review safety information.	20%
PO01	Create a purchase requisition, and then create a purchase order from the purchase requisition.	5%
PO02	Change the status of a purchase order to In Progress.	5%
PO03	Receive a purchase order line item.	5%
PO04	Close a purchase order.	5%
WF01	Create a work order and route the work order through the Workflow application.	12%
WF02	View a work order and route the work order through the Workflow application for approval.	12%
WF03	Issue an item on a work order and route the work order though the Workflow application.	12%
WF04	Add labor to a work order and route the work order though the Workflow application for completion.	12%
WF05	Route a work order though the Workflow application for closure.	12%

Step Three: Develop tests

The deployment team writes each use case into a test case. Each test case lists each step that is required to run the test. The following table provides an example of the test case for use case AS01, which searches for assets and then reviews safety information.

Table 28. Example test case to search for assets and review safety information

Transaction	Description	Expected result
AS01_01_D_Launch	Start Maximo Asset Management.	The Welcome to Maximo screen is shown.
AS01_02_D_Logon	Enter the user name ASSET0001 and the password maxasset. Click Sign In .	The Start Center is shown.
Begin loop for multiple work items.		
AS01_03_D_GoTo	Click Go To .	The Go To menu is shown.
AS01_04_D_LaunchAssets	Select Assets > Assets	The Assets application is shown.
AS01_05_D_EnterAsset Prefix	In the Asset field, enter CAC and press the Tab key.	The background for the Asset field changes to white. The cursor moves to the next field.

Table 28. Example test case to search for assets and review safety information (continued)

Transaction	Description	Expected result
AS01_06_D_FindAsset	Click the Filter Table icon.	A list is shown that lists all assets that have CAC in their names.
Loop up to 9 times to select a random page of data.		
AS01_07_D_NextPage	Click the Next Page icon	The next page of asset results is shown.
End loop for page data.		
AS01_08_D_SelectAsset	Select a random asset number.	The details for the selected asset are shown.
AS01_09_D_TabSafety	Select the Safety tab.	The Safety tab is shown.
AS01_10_D_ReturnToStartCenter	Click Start Center .	The Start Center is shown.
End loop for multiple work items.		
AS01_11_D_Logoff	Click Sign out .	The logout is completed. The Welcome to Maximo screen is shown.

Step Four: Define the test environment

In the initial planning stages, the deployment team discusses whether the cost of a test environment that is identical to the production environment is a justifiable expense. In the end, the deployment team decides that the risks of an inadequate test environment outweigh any potential cost savings. Therefore, the test environment at Company XYZ is an exact duplicate of the production environment.

In preparation for deployment, existing data from the system that Maximo Asset Management is scheduled to replace is migrated into the test environment. The migration of data ensures that the team is able to migrate existing data and also provides a realistic volume and structure of the data in the database that is then used for performance testing.

After the initial deployment into production, the team can test additional modifications in the test environment. The similarities between the test and production environments provide a high degree of confidence that similar results can be achieved when the additional modifications are moved to the production environment.

Step Five: Run tests

The deployment team can now record the test for the example test case and repeat the process to develop test cases for all the use cases. The deployment team uses a performance test tool to create the test cases. After all tests are recorded and debugged, the deployment team runs the tests.

To learn how the system operates at different load levels, the deployment team begins the test with 750 concurrent users. The user load is increased by an additional 50 users after a 30-minute interval. The increase of users is repeated until 950 concurrent users are on the system. The test is configured to log in one virtual user every 2000 milliseconds until the users at each load level are logged in.

This process is intended to eliminate the extra processing that is required to increase the load.

Step Six: Analyze test results

After the tests are run and the data is compiled, the deployment team extracts the response time and processor utilization results into a spreadsheet. The team can generate a summary chart to identify whether the performance criteria are met. The following chart shows an example of utilization results:

Response times and processor usage for 750-950 users




Figure 3. Example of performance test results for processor utilization

In the results chart, the average response times are under 2 seconds. The processor utilization on the database server remains below 80% at the target load of 950 concurrent users. However, the application server processor exceeds 80% utilization with a load of 850 concurrent users. Therefore, the performance test criteria were not met.

The deployment team must investigate to determine whether the excessive processor utilization issue can be resolved by tuning performance-related settings or changes to the automated workflows. The deployment team can also decide whether additional processor resources are required to meet the performance criteria in the production deployment.

Related tasks:

“Determining test objectives” on page 130

Performance tests can validate sizing estimates, ensure that your deployment meets your business or organization requirements, and address performance issues. You must identify the key business questions for your organization and develop tests that answer those questions.

“Developing use cases” on page 131

When you design performance tests, your goal is to develop few use cases that adequately test the functions that are most frequently used and most important.

“Developing test strategies” on page 131

The type of performance test that you use depends on your test objectives. Focus your test strategy on identifying the type of performance test that can provide data to measure the objectives.

“Defining test environments” on page 132

Your performance environment is composed of many different components, such as operating system, middleware, and deployment topology. Plan the specifics of each component level to create a test environment.

Chapter 6. Implementing security

You can implement security to manage which users can log in and which applications users can access. You can set up security privileges for users by group. You create the groups using the Security Groups application. For each group, you specify security settings, such as application privileges and restrictions. You grant security rights to users by assigning them membership in one or more groups. The combination of groups to which a user belongs determines the security privileges for the user.

Security Groups overview

You use the Security Groups application to set up and to manage security privileges for users. You can set up security privileges for users by group. You grant security rights to users by assigning them membership in one or more groups.

You can assign users to groups from both the Security Groups application and from the Users application:

- In the Security Groups application, you assign users to groups.
- In the Users application, you assign groups to users.

You can view the security privileges for a user using the **Security Profile** tab in the Users application.

Upon implementation, the Security Groups application has the following groups:

- **MAXDEFLTREG** - This group allows users to change their password if it expires. The group contains no other rights. When you create a user record, the user is placed in this default group. You can specify a different group to be the default using the **Security Controls** action. If you want new user security profiles to start with more rights, you can change the MAXDEFLTREG group to include these rights.
- **MAXADMIN** - This group provides enough access to add users and groups.
- **MAXREG** - This group allows users to self register. You can use the MAXREG group to initiate a workflow process by which an administrator is alerted to assign new users to the appropriate security groups.
- **MAXEVERYONE** - This group is used for global settings that apply to all users in the system.

When you delete users and security groups on the Lightweight Directory Access Protocol (LDAP), the users and security groups are not deleted in the system. This restriction is for audit purposes for clients in regulated industries.

Multisite implementation

If your company has multiple sites, you can create groups to reflect these sites. You can then combine the site groups with functional groups to create fine-grained sets of security privileges. For example, if you have sites in Toronto and Montreal, you can name two groups TORONTO and MONTREAL. You can then add groups to reflect functional units, such as finance, administration, maintenance, electrical, and so on.

Application server

If implementation uses an application server to authenticate with an external directory, some functions are performed in the directory and synchronized into the system. These functions include adding users (including self-registration), adding security groups, associating users with security groups, and managing passwords. In addition, when you delete users in the directory, those users are not automatically deleted in the system; you must manually delete them.

Security groups and access to sites and applications

Security access is based on security groups. You configure security groups to provide narrow access or broad access to applications, sites, and labor. You can also provide access to general ledger components, approval limits and tolerances.

Sites

The security architecture is designed to use sites as the first level of security for multisite implementations. You might consider using the following strategies when implementing security:

- If your implementation has only one site, then for each group, use the option to authorize the group to access all sites.
- If your implementation has multiple sites, create groups to represent each site, all sites, or a logical grouping of sites within a security group. For example, you can create a security group for site 1, and a security group for site 2 and site 3.
- Do not include any other privileges for the site groups.
- An independent security group has access rights and grants that cannot be combined with the rights and grants from other groups. If you select this option, you must grant that group access to at least one site and one application. You grant access unless the group is being used exclusively for system-level applications.

Applications, storerooms, labor, general ledger components, limits and tolerance, and restrictions

The following strategies apply to applications, storerooms, labor general ledger components, limits and tolerance, and restrictions:

- You can create groups that reflect these privileges. For example, if your company or facility has four storerooms, you can create separate groups for each storeroom and a fifth group for all storerooms. You can add those groups to the profile for a user, as appropriate.
- You can create functional groups that combine some of the privileges. For example, you can create three different maintenance groups. Each group would have different levels of privileges for any or all the properties in the tabs in the Security Groups application. This strategy is good for defining groups in a detailed manner. When you associate one group with a user, the group encompasses all or many of the privileges that you want the user to have.

Depending on how you want to implement security, you can also create groups that use a mixture of these two approaches.

Additionally, the following rules apply to when security privileges include access to applications:

- When you have access to a system-level application, any changes that you make in that application have a system-wide impact. For example, if you have access to the Currency application and add EURO as a currency, that currency is available for all organizations and sites.
- When you make a change in an organizational-level application, the change applies to all sites in the organization. For example, you are a user at site 1. You have access to the Chart of Accounts application, and make a structural change to an account. The change affects all sites within the organization to which site 1 belongs.
- Any changes that you make within a site-level application are limited to that site.
- The level of the application controls the amount of data that you can view. For example, site-level applications list data for specific sites, and organizational-level applications list data for all sites within an organization.

Types of security groups

To provide flexibility when you build the security infrastructure for your organization, you can choose from two types of security groups. There are independent security groups and groups that are not independent.

When you create a security group in the Security Groups application, you can specify whether the security group is independent or is not independent. If you create a security group that is not independent, the access rights and grants in the security group are combined with the rights and grants in similar security groups. The access rights and grants for an independent security group cannot be combined with the rights and grants from other security groups.

You can choose from the following options for combining the two types of security groups:

- Do not combine - results in all independent security groups
- Combine - results in security groups that are not independent
- Do not combine and combine - results in a combination of independent groups and groups that are not independent

Security process

The Security Groups application uses a two-step security process that consists of authentication and authorization.

Authentication is the process of validating the identity of a user. There are different authentication methods that share a common trait: authentication is always provided by a user ID and password. After a user is authenticated, authorization lets that user access various resources based on identity.

Authorization determines which modules and applications a user can access, which actions a user can perform, and which data a user can view, change, and delete. Authorization is provided by membership in one or more security groups.

Authentication of users

In the Security Groups application, you can set up authentication to validate the identity of a user. Authentication is the process of validating the identity of a user through a user ID and a password.

You can authenticate users through the following methods:

- You can use the application server and a Lightweight Directory Access Protocol (LDAP) server, either with or without Virtual Member Manager. LDAP is a set of protocols to access information directories.
- You can use a Web client server for native authentication.

Authenticating using LDAP

You can use application server security with an external authentication mechanism, such as LDAP, to authenticate users. The system uses application server security with an external authentication mechanism.

The system is built with Java 2 Platform, Enterprise Edition (J2EE) technology. This technology requires a commercial application server. The system uses WebSphere Application Server or WebLogic Server. By default, WebSphere Application Server security is enabled.

Authenticating using LDAP with Virtual Member Manager

You can authenticate users against LDAP using Windows Server Active Directory and Virtual Member Manager.

When you configure the application server to authenticate against an active directory, you create and manage users in the LDAP directory server. The Virtual Member Manager cron task updates the database when users, groups, and group membership are changed in the directory server. When users and groups are deleted from the active directory, they are not deleted from the database. This occurs because these records could be needed for auditing purposes.

You can also configure the system to populate person, user, and group information from the external directory. The system currently supports synchronization of information from Microsoft Active Directory. Synchronization with other directories is possible, but is not supported as a standard feature and can require programming to configure.

Both WebLogic Server and IBM WebSphere Application Server support authentication against Windows Server Active Directory.

Authenticating using a Web client server for native authentication

You can use the native authentication provided with the system to authenticate users and verify their identity and security authorizations.

When a user provides a login ID and password, the security functions validate whether the user ID and password are in the database. The user is granted access to applications, actions, and data based on the security groups with which their user ID is associated.

In addition, the security services perform the following actions at startup:

- Verify if the login ID is blocked or inactive.
- Authenticate the login ID and update password history.
- Establish the default insert site, organization, and person ID for the user.
- Establish the language, locale, time zone, and start center ID for the user.
- Route any workflow assignments to the inbox for the user.

Authorizations for security groups

You use the Security Groups application to grant authorizations to security groups.

You can grant the following authorizations to security groups:

- Start center and sites
- Read, insert, save, and delete access to applications and access to menu options
- Inventory storerooms and labor records
- Purchasing limits and tolerances, and general ledger components
- Conditional data restrictions for objects and attributes

If you implemented Lightweight Directory Access Protocol (LDAP), you create and manage security groups in the LDAP server (unless you changed the default behavior using the `useappserversecurity` property). To create and manage groups in the LDAP server, put users in security groups. The combination of security groups represents security profiles for users. Users acquire the authorizations and rights of the security groups to which they belong.

Related tasks:

“Authorizing application privileges for security groups” on page 164

According to your security needs, you can grant a group specific privileges within an application. These privileges include read, insert, save, and delete.

Application access for security groups:

In the Security Groups application, you can grant users access to specific applications to refine security measures. Users can have read, insert, save, and delete access to an application. The application access of a security group is linked to site access. You can give a security group access to all sites, access to specific sites, or no access to sites.

You can grant users specific options within an application. For example, you can grant managers the right to read work order histories, costs, and warranties, but not to insert work orders or service requests. You must configure each application for read access so that administrative users can select additional application access options.

All applications and their corresponding access options appear in the `SIGOPTION` table, which contains the following types of column information:

- Application option description
- Application option name
- Visible
- Also grants
- Also revokes
- Prerequisite

The visible setting indicates whether a user can select the option in the Security Groups application. If an option is not visible, the option is granted with another option. The standard system options that are not visible include clear, bookmark, next, previous, viewhist, and drilldown. For example, when you grant read access, the invisible options, clear, bookmark, next, previous, viewhist, and drilldown, are granted.

The values for the also grants, also revokes, and prerequisite access options indicate relationships between options. For example, if you grant insert access for an application, the also grants access option grants save access.

Standard access options are associated with prerequisite, also grants, and also revokes, as described in the following table.

Table 29. Standard options for prerequisite, also grants, and also revokes access

Standard options	Relationships between options
Standard prerequisite • Duplicate • Delete	Prerequisite • Insert • Save
Standard also grants • Insert • Read	Also grants • Save • Clear, bookmark, next, previous, viewhist, drilldown
Standard also revokes • Read • Save • Insert	Also revokes • All options • Insert, duplicate • Duplicate

The relationships in individual applications can sometimes vary. To view option access information for a specific application, use an SQL editor to search the SIGOPTION table.

Security groups and start centers:

Start centers are assigned to security groups. Through portlets, the start center allows quick access to the tools and key performance indicators that users access.

When users first log in, they see a start center based on a template for their security group. If users belong to more than one security group, they can see tabs representing a start center page for each security group.

As an administrator, you can grant users authorization to configure their start centers. You control the portlets that users can view and can configure. A start center can contain portlets for the following:

- Bulletin board
- Favorite applications
- Inbox
- Key performance indicator graph
- Key performance indicator list
- Quick insert
- Result set

Conditional security:

You can apply these conditions to security groups, to enforce security measures, for example, you can limit access to certain elements of the user interface to users with the right level of access.

In the condition library, you can define conditions, either as expressions or as custom class files. You can control access to applications and to controls in applications by applying conditions to security groups. The conditions are applied to signature options (SIGOPTION) that are then granted to the security group. Conditional access is granted in the security groups application. If a user is in multiple security groups, the highest level of access is granted when the security groups are joined.

The following are some examples of the types of conditional access that you can set:

- Give read-only access to the information that is shown in a field.
- Give read/write access to the information that is shown in a field.
- Give a user group read-only access to a specific field in an application.
- Give all members of a user group read/write access to an application.
- Hide a field or tab in an application from certain users.
- Grant access to application options in the **Select Action** menu for a security group.
- Configure any property in a control for a group, such as making a control hidden, masked, read-only, or required.
- Configure other properties, such as color, label, and application link, to be different for on the user group that is accessing the UI.
- Show or hide a data attribute globally or for a security group.

Conditional security and data restrictions:

In the Security Groups application, you can use data restrictions to meet conditional security requirements for users. You can set restrictions on which records a group can access within the larger set of records.

You can use data restrictions to limit the data to hide records or to make records read-only. At the attribute level, you can create data restrictions to make records hidden, read-only, or required. Because these data restrictions exist at the data-level, the restrictions apply to any user interface element or application that uses an object or attribute.

Data restrictions provide the following ways to configure access to data for groups of users:

- You can make an entire object or an entire object within the context of an application hidden or read-only, either conditionally or unconditionally for the entire system or for a security group.
- You can associate an object or object and application with a condition to qualify the data to be returned. Only data that meets the condition is fetched from the database. This differs from data that is fetched from the database but is hidden in a certain condition. Qualified data restrictions are applied only to top-level objects in lookups and in dialogs that are configured to allow them.
- You can set data restrictions for attributes within objects, either with or without an application specified. In these restrictions, you can make the attribute hidden, required, or read-only, either conditionally or unconditionally, for the entire system or for a security group. At run time, within the applications, controls bound to restricted objects or attributes can change their display as a user scrolls through records.
- You can set collection restrictions to control the collections of assets, locations, and configuration items that a group can access.

- Data restrictions always supersede application configurations in the Application Designer application. For example, if an attribute has a data restriction that makes it read-only, the Application Designer application can never make that attribute editable. The hierarchy is database configuration, data restriction, and then Application Designer application.
- Configurations that you create with data restrictions apply wherever an attribute is used, while Application Designer configurations do not. For example, you want to restrict access to a field that appears in the header section of multiple tabs. If you put a data restriction on the attribute, all the fields inherit the restriction. If you configure the same restriction in the Application Designer application, you must apply the same configuration to each field on each tab.
- Application Designer configurations are always for one application. Configurations that use data restrictions can apply to all applications that use the object or attribute or to one specific application.

If you create a data restriction on an object, that restriction does not apply to views of that object. For the restriction to apply to all views of the object, you create a separate restriction for the view.

When you grant a user access to an application, the user has access to all the data elements per the business logic of that application.

Group data restrictions

In the Security Groups application, you can set restrictions using a condition that defines which records a group can access. If a user is in multiple groups, and one or more of those groups has data restrictions, the data restrictions behave in certain ways: qualified data restrictions are ORed together and other data restrictions are ANDed together.

However, if one of the groups has application access, then different rules apply. If a user belongs to a group with read access and also has access to a siteorg, then data restrictions are considered. If not, then data restrictions are ignored.

Global data restrictions

You use the **Global Data Restrictions** action to set restrictions that use a condition that defines which records can be accessed in the system. To create expressions for these conditions, use the Conditional Expression Manager application.

Security profiles

A security profile is the list of rights a user derives from the security groups to which the user belongs. The user rights define the ability of a user to access asset management system applications and to perform specific application functions. You can view a user security profile list in the **Security Profile** tab in the Users application.

Before you can view a user security profile list, you must first create a user in the Users application. You must then assign the user to a security group and allocate group rights in the Security Groups application. When you assign a user to more than one group, the user inherits the security rights from the assigned groups.

You can use the Security Groups application to define user rights according to the following asset management system security components:

- Site

- Application
- Storeroom
- Labor
- GL Component
- Limit and tolerance
- Data restriction

By default, the asset management system includes the following user rights:

- The ability to change and a password at login, even when a password expires
- The ability to access the asset management system start center

If your implementation uses a Lightweight Directory Access Protocol (LDAP) directory server to authenticate an external directory, do not use the asset management system security settings to perform the following functions:

- Self registration
- Password and password hint changes

Security profile of an organization with two security groups - example

A security profile lists the access that a user has after all the security groups are combined. This example shows an approach to building a security profile for one organization with two security groups.

In the following example, an organization has two security groups: the worker security group and the management security group. The security profile is the result of combining the groups to restrict access to applications, and limits and tolerances.

For example, the security profile for the worker security group does not provide access to purchase requisitions and financials, and does not have a purchase limit.

Example: Single organization with security groups that provide sufficient application, site and storeroom access and privileges for all users in XYZ company.




Figure 4. Organization with two security groups

Example: Multi-organizational implementation that uses independent, non-independent and the 'Everyone' security groups to provide the user with read-only access to certain applications for certain sites, more robust access to applications in other sites and conditional access to options. This example illustrates that the 'Everyone' group combines with both independent and non-independent groups. It also illustrates that a user with 'full' access to an option in one group and conditional access to the same option in the 'Everyone' group will get 'full' access to the option.





Figure 5. Multi-organizational implementation

Login tracking

You use login tracking to specify the number of allowed login attempts. You can track the number of login attempts and view the current login status for a user. You can also block further attempted logins by a user who exceeds the specified number.

You use the **Security Controls** action to enable login tracking and to specify the maximum number of unsuccessful logins. Login tracking is required if you use electronic signature.

When login tracking is enabled, all successful and unsuccessful login attempts are tracked. After each successful login, the maximum number of chances to log in is

reset. When users reach the maximum number, their status is changed to blocked. They are prevented from logging in until an administrator changes their status.

If implementation uses an application server to authenticate with a directory, some functions are performed in the directory and synchronized into the system. These functions include adding users (including self-registration) and managing passwords.

Encryption and security

The data types Crypto, and CryptoX are used to encrypt passwords and other types of confidential information. The Java Cryptography Extension (JCE) is used to perform encryption.

JCE can use variables to transform the input data into encrypted data. By default, the DESede encryption algorithm and its defaults are used for the other values. Crypto and CryptoX use the DESede encryption algorithm.

The following table describes the data types Crypto and CryptoX.

Table 30. Crypto and CryptoX data types

Data type	Data stored	Algorithm
CryptoX	User passwords	• One-way encryption • Stores password in encrypted format (cannot be decrypted or displayed) • Internally, the encrypted version is used
Crypto	Information that you want to decrypt or display	• Two-way encryption • Information can be decrypted and displayed to users

You can configure the encryption settings in the `maximo.properties` file. You can configure the encryption data types to be consistent with industry guidelines and government guidelines. You can also configure encryption to make your system more secure: key, mode, padding, and spec. The following table describes the encryption settings.

Table 31. Encryption settings

Encryption property	Settings for JCE and DESede
<code>mxe.security.crypto.key</code> <code>mxe.security.cryptox.key</code>	Length must be a multiple of 24
<code>mxe.security.crypto.mode</code> <code>mxe.security.cryptox.mode</code>	• CBC - Cipher Block Chaining Mode • CFB - Cipher Feedback Mode • ECB - Electronic Codebook Mode • OFB - Output Feedback Mode • PCBC - Propagating Cipher Block Chaining
<code>mxe.security.crypto.padding</code> <code>mxe.security.cryptox.padding</code>	• NoPadding • PKCS5Padding

Table 31. Encryption settings (continued)

Encryption property	Settings for JCE and DESede
<code>mxe.security.crypto.spec</code>	Length must be a multiple of 8
<code>mxe.security.cryptox.spec</code>	

Hacking and denial-of-service attacks

Malicious users can attempt to breach security or affect system performance by attacking the login, forgotten password, and self-registration processes. You can configure system properties and security settings to help prevent these kinds of hacking and denial-of-service attacks.

IP address blocking

Depending on your security settings, IP addresses are blocked when an attack is detected. You can view, add, and delete blocked IP addresses in the Manage Blocked IP Addresses window of the Users application.

Security can be configured to block an IP address when too many login, forgotten password, or self-registration attempts are made from the same address. For any blocking to occur, the **`mxe.sec.IPblock`** property must be set.

In addition, if the **`mxe.sec.IPblock.MatchBoth`** property is set, an IP address is blocked only if both the client host and the client address of the incoming request match the values in the LONGINBLOCK table.

Failed logins and forgotten password attempts

If the number of failed logins or forgotten password attempts from the same IP address exceeds the value of the **`mxe.sec.IPblock.num`** property in the time that is specified by the **`mxe.sec.IPblock.sec`** property, the IP address is blocked.

Furthermore, if the number of concurrent forgotten password attempts exceeds the value of the **`mxe.sec.forgotpassword.maxsets`** property, an error occurs and the requesting IP address is blocked.

The number of failed logins is tracked by the reported number of web browser sessions. However, do not use the **`mxe.sec.IPblock.num`** property to try to control the number of user sessions or windows. Different web browsers report sessions differently, depending on the use of tabs, the browser version, and the operating system. Therefore, the number of browser sessions might not match the number of browser windows. Set the **`mxe.sec.IPblock.num`** property only for purposes of blocking intrusion attempts.

If the number of successive forgotten password attempts for a user exceeds the value specified in the Security Controls window of the Users application and the Security Groups application, the status of the user that is associated with the email address is set to BLOCKED.

Self-registration attempts

If the number of concurrent self-registration attempts exceeds the value of the **`mxe.sec.addusers.maxsets`** property, an error occurs and the requesting IP address is blocked.

Automatic creation of user records authenticated by LDAP

The system can create user records that pass Lightweight Directory Access Protocol (LDAP) authentication. This function allows you to create users with basic privileges and then reassign the users to the appropriate group.

You use the Properties application to enable this function. In this application, set the value of `mxe.Allow LDAPUsers` from the default value of 0 to 1.

Combination of security groups

In the Security Groups application, you can combine security groups to manage the security infrastructure within or across organizations.

When you combine security groups, the following rules apply:

- You cannot combine the privileges of independent security groups.
- You can combine the privileges of security groups that are not independent.
- When you combine privileges, the highest privileges prevail. If a user belongs to multiple security groups that define the same privilege at different levels, the user has the highest privilege. For example, security group A has a purchase order limit of \$5,000. Security group B has a purchase order limit of \$10,000. A user who is a member of both security groups has a purchasing limit of \$10,000.
- When you combine a security group that has data restrictions, the restrictions are added to the security profile for the user. This action can reduce the access rights that were otherwise granted by the combined security groups.
- Using the Security Controls **action** in the Security Groups application or Users application, you can specify the group for all users, the **MAXEVERYONE** group. The **MAXEVERYONE** group always combines, even if the group is specified as independent.

Combining security in a multiple site environment

Combining privileges is useful when you have multiple sites. Typically, you set up security groups that only define site access. You set up other security groups to define application privileges, purchasing approval limits, and so on. For example, your organization has three sites, site 1, site 2, and site 3. You have a user for whom you created a security profile that includes site 1 and associated privileges. You want the user to have the same privileges at site 2, therefore, you add site 2 to the profile for the user.

You can also define some security groups as independent, so that when you combine security groups, a user has a set of privileges at one site and a different set of privileges at another site.

Combination of security groups - rules for data restrictions

When you combine independent or non-independent security groups, you use restrictions to specify the records that are visible to members of a security group.

When you combine security groups and use restrictions, the following rules apply:

- If a user is a member of multiple groups that are not independent and one security group has a restricted level of access, the user is granted the highest privileges across the security groups. For example, take two security groups that are not independent: the Managers security group and the Maintenance security group. The user has access to pay rate information in the Managers security group, but does not have access to the information in the Maintenance security group.

group. When the two security groups are combined, the user has access to pay rate information in the Maintenance group.

- Data restrictions always combine across security groups by using the OR operator regardless of whether the groups are marked as independent. For example, take two security groups: one security group contains a READONLY data restriction condition ":orgid [equals character] 'EAGLENA'". The second security group contains a READONLY data restriction condition ":orgid [equals character] 'EAGLEUK'". Regardless of whether one, both, or neither security group is marked as independent, the restrictions combine to make the object or attribute read only if the ORGID is EAGLENA_OR_ EAGLEUK.

Therefore, you must consider the conditions that you apply to data restrictions carefully.

Combination of security groups - rules for application authorization

In the Security Groups application, you combine the independent security groups or the security groups that are not independent to generate a security profile for a user.

When you combine security groups, the following rules apply to application authorization:

- The application authorizations specified for an independent security group apply exclusively to the sites or organizations associated with that security group.
- The application authorizations specified for all security groups that are not independent apply to all sites specified for those security groups.
- When you do not specify a site, access to the Users application and to the Security Groups application is site independent. To change site administration, you can specify a site for the security group that grants access to these applications. For example, if a security group has access to these applications at site 1, a user who is logged in can only change information for users who are associated with site 1.
- The available options in the **Select Action** menu and **Go To** menu depend on the options and applications that you granted a user, regardless of the site or organization.
- You can grant a user access to the **Change Status** action based on sites.
- If the security profile for a user has application authorizations but no sites, the user can access the applications. The user cannot view or insert records, except for the Users application and Security Groups applications.
- You can define conditional access to applications. For example, a user is a member of two security groups, and one security group has conditional access and the other security group has unconditional access. The unconditional access overrides the conditional access, and the user has unconditional access.

The accumulation of all unique application authorization records across security groups becomes the access list of application authorizations in the security profile of a user.

Combination of security groups - rules for approval limits and tolerances

In the Security Groups application, you combine the independent security groups or the security groups that are not independent to generate a security profile.

When you combine security groups, the following rules affect approval limits and tolerances:

- The limits and tolerances that you authorize for a security group are at the organizational level. Users inherit authorizations for only those sites to which they have access.
- If there are two different approval limit values for the same limit, the higher value is applied to the security profile for the user.
- If there are two different values for the same tolerance type, the higher value is applied to the security profile for the user.
- If there are two different values for the same tolerance type, but the security groups that grant the tolerance amount have sites in different organizations, the higher value for sites within the same organization is applied to the security profile for the user.
- If a user has access to two different organizations with different limits and tolerances in each organization, the user inherits the limits and tolerances for each site to which the user has access in each organization.

The security profile for a user lists the specified approval limits and tolerances.

Combination of security groups - rules for authorization of general ledger components

In the Security Groups application, you combine the independent security groups or the security groups that are not independent to generate a security profile.

When you combine security groups, the following rules apply to the authorization of general ledger components:

- If any of the security groups to which a user belongs grants authorization to change all general ledger components or specific general ledger components, the security profile for the user reflects the maximum amount of general ledger component authorization.
- If you do not authorize a security group to change all general components, and you do not authorize individual components for the security group, a user cannot change general ledger components.
- The general ledger component authorizations specified for all security groups that are not independent apply to all of the applications, sites, and organizations for those security groups for the user.
- The general ledger component authorizations specified for an independent security group apply exclusively to the applications, sites, and the organizations associated with that security group.

The security profile of a user lists the authorizations for general ledger components.

Related tasks:

“Authorizing security group access to general ledger components” on page 166
As a security measure, you can authorize which general ledger components a group can access.

Combination of security groups - rules for labor authorization

In the Security Groups application, you combine the independent security groups or the security groups that are not independent to generate a security profile.

When you combine security groups, the following rules apply to labor authorization:

- All labor in an organization.
- All labor in the same crew as the user.
- All labor in the same person group as the user.
- All labor that the user supervises.
- Only the labor records for the user.
- Individual labor records that are listed in the table window.

The security profile of a user lists the authorized labor.

Combination of security groups - rules for site authorization

In the Security Groups application, you combine the independent security groups or the security groups that are not independent to generate a security profile.

In the Security Groups application, you can give members of a security group access to all sites. If a user does not have access to all sites using group membership, the sites to which a user has access through group membership is tracked. A cumulative list of sites is included in the security profile for the user.

Combination of security groups - rules for storeroom authorization

In the Security Groups application, you combine the independent security groups or the security groups that are not independent to generate a security profile.

When you combine security groups, the following rules apply to storeroom authorization:

- A user must have access to both a storeroom and the site for the storeroom before the storeroom authorization is added to the security profile for the user.
- The storeroom authorizations that are specified for an independent security group apply exclusively to the sites associated with that security group.
- In the Security Groups application, you can authorize a security group to make transactions with storerooms. The storeroom authorizations specified for all dependent security groups apply to all sites that are specified for all dependent security groups.
- If you do not authorize a security group access to all storerooms or access to individual storerooms at a site, a user who is a member of that security group does have access to storerooms at that site.
- If any of the security groups to which a user belongs grants access to all or specific storerooms at a given site, then the security profile for the user reflects the maximum amount of storeroom access.
- You can give a user access to the Users application and Security Groups application, but not grant that user access to any storeroom records. In this scenario, the user can use these applications to authorize access to all storerooms, but cannot add specific storerooms records.

The security profile of a user lists the authorized storerooms.

Application server security

You use the Security Groups application to configure application server security.

Application server security - properties for user and group management

After you enable application server security, you can configure the properties for group and user management to define whether the directory owns group creation, or user creation and management.

By default, the `mxe.LDAPGroupMgmt` property is set to 1. This setting indicates that the directory owns group creation and group management. When you change the value to 0, the system owns group creation and group management. This setting enables the following functions:

- Create security groups
- Associate users with security groups

By default, the `mxe.LDAPUserMgmt` property is set to 1. This setting indicates that the directory owns user creation and user management. When you change the value to 0, the system owns user creation and management. This setting enables the system to use Lightweight Directory Access Protocol (LDAP) for user authentication without having to synchronize user information. The following table lists the functions that are enabled and disabled with the 0 setting.

Table 32. Enabled and disabled functions when the system owns user creation and user management

Function	Enabled
Add and delete security groups	No
Change security groups	Yes
Manage user and group relationships	No
Add and delete users	Yes
Change users (other than password)	Yes
User self-registration	No
Change password	No

The user ID records that are created in the directory and in the system must be identical for this setting to function correctly.

Related tasks:

“Configuring the system to use application server security” on page 174

You must configure the system to use application server security. However, if application server security was configured through the installer, you do not have to complete this procedure.

“Configuring WebLogic Server for LDAP security” on page 177

To implement configuration to use application server security, you configure your application server to use Lightweight Directory Access Protocol (LDAP).

“Configuring WebSphere Application Server for LDAP security” on page 175

To implement configuration to use application server security, you configure your application server, configure to use Lightweight Directory Access Protocol (LDAP).

.

Security roles for the application server

When you enable application server security, you can use roles to manage security.

The following table describes the security roles that you can use for IBM WebSphere Application Server and WebLogic Server.

Table 33. Security roles for WebSphere Application Server and WebLogic Server

Application server	Security role	Description
WebSphere Application Server	Supports scope roles only.	• Map the scoped roles to individual non-nested groups. • Might not authenticate users in nested groups on certain LDAP servers. Please refer to WebSphere Application Server documentation for details.
WebLogic Server	Supports both global and scoped roles. By default, uses scoped roles. Use the administration console to change to global roles.	• Global - applies to all resources within a security realm (the entire server domain). • Scoped - applies to a specific instance of a resource that is deployed in a security realm.

For more information about securing resources for WebLogic Server, go to the WebLogic Server Web site and search for securing resources. For more information about types of resources for WebLogic Server, go to the WebLogic Server site and search for types of resources.

For more information about WebSphere Application Server, go to the IBM Information Center for WebSphere Application Server Network Deployment, and search for security roles.

Single sign-on environment for application server security

When you enable application server security, you can use a single sign-on environment. This environment enables a user to provide one name and password to access multiple applications.

When users authenticate with the server in a single sign-on environment, the users are authenticated to access all of the applications to which they have been given rights on the server. This authentication eliminates the need to provide multiple passwords when users switch applications.

Configuration for a single sign-on system depends on your implementation. For more information about how to configure your single sign-on environment so that the system can participate, see the documentation for your single sign-on platform and your application server. Both the WebLogic Server and WebSphere Application Server support a single sign-on environment.

LDAP and application security servers

The application server is configured to authenticate against a Lightweight Directory Access Protocol (LDAP) server user registry. The system supports integration with Microsoft Active Directory or IBM Tivoli Directory Server LDAP server. You can move LDAP server data into product database tables. The application server that you use determines which directory provides support.

You add users and delete users and groups from the LDAP server, but the system provides the authorization. By default, the property `mxe.LDAPGroupMgmt` is set so that group creation and group membership is managed by the directory server.

You can configure all application-specific authorization rules for users and groups using the security module applications. In the system, you disable password information in the start center, the Change Password application, self-registration, and the Users application.

LDAP server users and groups are moved into product database tables to identify users as system users, and to provide user with details in system applications.

Users and groups that are deleted from the LDAP server are not deleted from database tables; audits might be conducted for users or groups.

If user accounts are disabled from the LDAP server, the application server has to expire the users cached information.

Before users can access the system, the application server authentication must be passed. Application servers use role records to identify users and groups that have access to the system. All roles that were configured in an application are mapped to users or groups using application server-specific deployment descriptors or application server-provided administrative tools.

LDAP data synchronization

Data synchronization keeps system data current with data in the Lightweight Directory Access Protocol (LDAP) directory server. Synchronized data moves only from the LDAP directory server to the system.

To synchronize data, you must create users and groups in the LDAP directory server.

The data synchronization process depends on which application server you use:

- If you use WebSphere Application Server, data synchronization is governed by the federated repositories which are managed by Virtual Member Manager (VMM) and the VMMSYNC cron task.
- If you use WebLogic Server, data synchronization is managed by the LDAPSYNC cron task.

The LDAPSYNC cron task supports incremental synchronization. The VMMSYNC cron task supports full synchronization by default, but it can be configured to support the incremental synchronization of users and groups. If you want to synchronize users and groups incrementally, you must set the `ChangePolling` cron task parameter to 1 or true. Incremental synchronization can help with performance and planning.

Changes that you make in the LDAP directory server are not reflected in the system until you synchronize the data. Even after you synchronize the data, there are some considerations to be aware of:

- When you synchronize users and groups in the LDAP directory server with the system, the users and groups become users and security groups in the system.
- Groups or users that you rename in the LDAP directory server are not renamed in the system. Instead of renaming, delete the group or user and then create a group or user.

- If you disable a user account in the LDAP directory server, that user account is not disabled in the system.

Synchronization of cron task parameters for application server security

You must configure the parameters for the synchronization cron tasks LDAPSYNC and VMMSYNC. The LDAPSYNC and VMMSYNC cron tasks are required for security in the application server.

The table describes the parameters that you configure for the synchronization cron tasks, LDAPSYNC and VMMSYNC.

Table 34. Parameters for synchronization cron tasks

Parameter	Description
LDAPSYNC	
Credential	LDAP credentials
GroupMapping	The GROUP XML that the LDAP task uses
Host	LDAP connection host
Port	LDAP connection port
Principal	LDAP principal
SSLEnabled	LDAP connection secure sockets layer (SSL) enabled
SynchAdapter	LDAP synchronization adapter
SynchParameter	Parameter name, value pairs are delimited by a comma
UserMapping	The USER XML that the LDAP task uses
VMMSYNC	
ChangePolling	Virtual Member Manager (VMM) parameter for incremental synchronization
Credential	VMM admin credentials
GroupMapping	The USER XML that the VMM task uses
GroupSearchAttribute	VMM search attribute to query group records
Principal	VMM admin principal
SynchAdapter	VMM synchronization adapter
UserMapping	The USER XML that the VMM task uses
UserSearchAttribute	VMM search attribute to query user records

The LDAP directory server maintains an attribute list for each user or group. Each attribute has an associated data type, which you can query the server to see. The LDAPSYNC cron task and VMMSYNC cron task only support string or character data retrieval from the LDAP directory server.

The data mappings in the LDAPSYNC cron task and VMMSYNC cron task parameters map LDAP attributes to system table columns. For the LDAPSYNC cron task to create a database record, all of the required columns must contain data. If all of the required column data cannot be obtained from the LDAP directory server, you must specify default values. To specify default values for

columns, the value must be enclosed inside brackets; for example, {ABC} specifies the value ABC in the column. The value is case-sensitive.

The synchronization task also supports special substitute values to generate unique IDs and system dates. To generate a unique ID for a column, use the notation {:uniqueid}; to generate system date, use the notation {:sysdate}.

Working with security groups

In the Security Groups application, specify security restrictions and privileges for groups of users.

Adding security groups

Security privileges control user access to modules, applications, menu options, and data. All security access is based on security groups.

About this task

If implementation uses an application server to authenticate with a directory, groups might be created in the directory and synchronized into the system.

Users can specify a default application for their user profile. If users do not specify a default application and default applications are specified for security groups, the default application for security groups displays when users log on to the system. If users are assigned to multiple security groups where default applications are specified, the system opens the Start Center. If no default settings are specified, the system default application displays.

Procedure

1. In the Security Groups application, click **New Group**.
2. Specify a name for the group.
3. Optional: Provide a group description and a name of the start center that displays when a user in the group logs in. Users also can customize their start centers or choose a default start center when they belong to groups with different start centers.
4. Optional: If you do not want rights combined, select the **Independent of Other Groups** check box. By default, rights are merged when groups that include different sites are combined.
5. Optional: Specify a default application for the security group.
6. Save the group.

Results

Privileges or restrictions are not defined for the new security group.

What to do next

You can use functions in the Security Groups application to define the security for the group. You can add users in the Security Groups application or the Users application.

Related concepts:

“Security groups and access to sites and applications” on page 140

Security access is based on security groups. You configure security groups to provide narrow access or broad access to applications, sites, and labor. You can also provide access to general ledger components, approval limits and tolerances.

Assigning start centers for security groups

Through portlets, the start center provides quick access to the tools and key performance indicators that users typically access. Administrators can assign the default start center that users see when they access the system.

Procedure

1. In the Security Groups application, select the group whose start center you want to assign.
2. Specify the name of the start center that displays when a user in the group logs in.
3. Save your changes.

Related concepts:

“Security groups and access to sites and applications” on page 140

Security access is based on security groups. You configure security groups to provide narrow access or broad access to applications, sites, and labor. You can also provide access to general ledger components, approval limits and tolerances.

Assigning sites to security groups

Sites are part of the security architecture. You assign security groups to provide access to sites.

About this task

You can only add sites to which you have access. You can only add inactive sites to a group if your user record is authorized for inactive sites.

Procedure

1. In the Security Groups application, select the group for which you want to assign site access.
2. Click the **Sites** tab.
3. Select one of these options:
 - To authorize the group to have access to all sites, select the **Authorize Group for All Sites** check box. If you select this option, you cannot add rows.
 - To authorize the group to have access to individual sites, click **New Row** and specify the name of the site. After you select a site, the remaining fields are populated.
4. Save your changes.

Related concepts:

“Security groups and access to sites and applications” on page 140

Security access is based on security groups. You configure security groups to provide narrow access or broad access to applications, sites, and labor. You can also provide access to general ledger components, approval limits and tolerances.

Adding users to security groups

You grant users security rights by assigning them membership in one or more groups. The combination of groups to which users belong determines security privileges.

Before you begin

To add users to an existing security group, you must be authorized to reassign users to that group. Use the **Authorize Group Reassignment** action to grant authorization. If you created a security group, you are automatically authorized to reassign users to that security group.

About this task

If implementation uses an application server to authenticate with a directory, you can add users to security groups in the directory or in the system, depending on your settings.

Procedure

1. In the Security Groups application, select the group to which you want to add a user.
2. Click the **Users** tab.
3. Click **New Row**.
4. Select a user. The associated user information populates the other fields.
5. Save your changes.

Related concepts:

“Security groups and access to sites and applications” on page 140

Security access is based on security groups. You configure security groups to provide narrow access or broad access to applications, sites, and labor. You can also provide access to general ledger components, approval limits and tolerances.

Authorizing security group reassignments for users:

For security purposes, you can specify the users who are authorized to add or to remove users from a security group.

Procedure

1. In the Security Groups application, select the group for which you want to authorize group reassignment.
2. From the **Select Action** menu, select **Authorize Group Reassignment**.
3. In the Authorize Group Reassignment window, click **New Row**.
4. To authorize users, click **Select Users**.
5. Click **OK**.

Related concepts:

“Authentication of users” on page 141

In the Security Groups application, you can set up authentication to validate the identity of a user. Authentication is the process of validating the identity of a user through a user ID and a password.

Setting user defaults:

You use the Security Controls action to specify the defaults for user records. You can access the Security Controls action from either the Security Groups application or the Users application.

About this task

You can specify the following defaults for user records:

- Default security group for new users - New users are automatically assigned to a security group. The default group defines their security permissions until they are assigned to additional groups. The default group is MAXDEFLTREG. The permissions for this group are limited to access to the Start Center. Users can change their own passwords.
- Default status for new user records - The default status is NEWREG. The NEWREG status allows you to search for new user records. You can also route records into a workflow process.
- Group for all users - The default security group for global permissions is MAXEVERYONE. When you select **Group for All Users** in Security Controls, the user is added to MAXEVERYONE.
- Electronic signature dialog - When you select **Display User ID in the Electronic Signature Dialog** in Security Controls, the system displays the user ID in the window, and prompts the user to enter a password.

Your implementation might use an application server to authenticate with an external directory by means of the Lightweight Directory Access Protocol (LDAP). In this case, you do not use the system to perform some functions. These functions include:

- Self registration - This function is not supported in conjunction with an external directory.
- Setting or changing passwords and password hints - All password-related functions are managed by the directory.

By default, when you use an application server for authentication, the directory manages user and group creation. You can set properties to allow user and group creation to be performed directly in the system. The settings of these properties result in certain features being enabled or disabled in the system.

Procedure

1. From the **Select Action** menu, select **Security Controls**.
2. In the User Defaults section, specify the following defaults:
 - a. In the **Default Group for New Users** field, type the name of the group, or click **Detail Menu**.
 - b. In the **Initial Self-Registered User Status** field, type a user status.
 - c. In the **Group for All Users** field, specify the group for global permissions. The default is MAXEVERYONE.
 - d. Select the **Display User ID in the Electronic Signature Dialog** check box to display the user ID in the window when the system prompts users to enter their passwords. If you implement electronic signatures, you must enable login tracking.
3. Click **OK**.

Granting authorization privileges to security groups

You can grant authorization privileges for security groups, such as access to applications, to storerooms, and to labor information.

Granting administrative login authorization for database configuration

To configure the database in administration mode, you must have administrative login authorization. Authorization is granted by security group.

Procedure

1. In the Security Groups application, select a security group and then click the **Applications** tab.
2. In the Applications section, search for Start Center.
3. In the Options for Start Center section, select **Can Log In During Admin Mode** and click **Grant Listed Options for This Application**.

Related concepts:

“Authentication of users” on page 141

In the Security Groups application, you can set up authentication to validate the identity of a user. Authentication is the process of validating the identity of a user through a user ID and a password.

Authorizing application privileges for security groups

According to your security needs, you can grant a group specific privileges within an application. These privileges include read, insert, save, and delete.

Procedure

1. In the Security Groups application, select the relevant group.
2. Click the **Applications** tab.
3. In the Applications table window, select the application. The options for the selected application are listed in the Options table window.
4. Select one of the following options:
 - To grant the privileges to all listed applications, click **Grant Listed Applications**.
 - To remove privileges to all listed applications, click **Revoke Listed Applications**.
5. Select one of the following privileges:

Option	Description
Read	This privilege allows users to access the application and to view records. You must select this option before you can select any others.
Insert	This privilege allows users to create records. If you select this privilege, the save privilege is automatically selected.
Save	This privilege allows users to save changes to records.
Delete	This privilege allows users to delete specific records. Before a record can be deleted, internal checks are executed to prevent deletion of records containing information required by other records. If you select the delete privilege, you must also select the save privilege.
All Above	This privilege grants or revokes read, insert, save, and delete privileges for all listed applications.

The options that you selected are listed in the Options table window.

6. In the Options table window, select one of the following options:
 - To grant access to specific options, select the relevant check boxes.

- To grant access to all options, click **Grant Listed Options for This Application**.
7. Optional: In the Options table window, select a conditional expression that conditionally grants the group the privileges for that option.
 8. Save your changes.
 9. For your changes to take affect, log out and then log in again.

Related concepts:

“Authorizations for security groups” on page 143

You use the Security Groups application to grant authorizations to security groups.

“Application access for security groups” on page 143

In the Security Groups application, you can grant users access to specific applications to refine security measures. Users can have read, insert, save, and delete access to an application. The application access of a security group is linked to site access. You can give a security group access to all sites, access to specific sites, or no access to sites.

Authorizing access to storerooms for security groups

Storeroom information is used in transactions that affect inventory items and balances. As a security measure, you can authorize a security group to make transactions with specific storerooms.

About this task

You can only add storerooms that you are authorized to access.

Procedure

1. In the Security Groups application, select the group for which you want to authorize storeroom transactions.
2. Click the **Storerooms** tab.
3. Optional: Authorize the group to have access to all storerooms. If you select this option, you cannot create individual storerooms.
4. Optional: To authorize the group to have access to individual storerooms, perform the following steps:
 - a. Click **New Row**.
 - b. Specify a site and a storeroom. When you specify a storeroom value before you specify a site value, you choose from all storerooms, including those with the same name that are in different sites. After you select a storeroom, the **Site** field is automatically populated.
5. Save your changes.

Related concepts:

“Combination of security groups” on page 152

In the Security Groups application, you can combine security groups to manage the security infrastructure within or across organizations.

“Combination of security groups - rules for storeroom authorization” on page 155

In the Security Groups application, you combine the independent security groups or the security groups that are not independent to generate a security profile.

Authorizing access to labor information for security groups

As a security measure, you can authorize a security group to access labor information, including records that contain labor information.

Procedure

1. In the Security Groups application, select the group for which you want to authorize labor information.
2. Click the **Labor** tab.
3. Select one of the following labor authorization options:

Option	Description
Authorize Group for All Labor	• Select this option for the group to have access to all labor. • If you select this option, the other labor-related selections are read-only.
• Authorize Group for Labor in Their Same Crew • Authorize Group for Labor in Their Same Person Group • Authorize Group for Labor They Supervise • Authorize Group for Their Own Labor	• Select one or more of these options for the group to have access to a limited set of records. • To select these options, ensure the Authorize Group for All Labor check box is clear.
New Row • Organization • Labor	• Select this option for the group to have access to individual labor records. • If you specify a value for labor first, the Select Value window lists labor from all organizations, including those of the same name in different organizations. After you select the value for labor, the Organization field is automatically populated. • You can only enter labor records that your user ID allows you to access.

- For the group to have access to all labor information, select the **Authorize Group for All Labor** check box. If you select this option, the other labor-related selections are read-only.
 - For the group to have access to a limited set of records, make sure the **Authorize Group for All Labor** check box is clear. Select options to authorize labor in the same crew, labor in the same person group, labor they supervise, and their own labor.
 - For the group to have access to individual labor records, click **New Row** and specify values for organization and labor.
4. Save your changes.

Related concepts:

“Combination of security groups” on page 152

In the Security Groups application, you can combine security groups to manage the security infrastructure within or across organizations.

“Combination of security groups - rules for labor authorization” on page 154

In the Security Groups application, you combine the independent security groups or the security groups that are not independent to generate a security profile.

Authorizing security group access to general ledger components

As a security measure, you can authorize which general ledger components a group can access.

Procedure

1. In the Security Groups application, select the group for which you want to grant general ledger component access.
2. Click the **GL Components** tab.
3. Optional: Authorize the group to access all general ledger components. If you select this setting, the settings in the GL Components table window are read-only.
4. Optional: Authorize the group to access individual general ledger components.
5. Save your changes.

Related concepts:

“Combination of security groups” on page 152

In the Security Groups application, you can combine security groups to manage the security infrastructure within or across organizations.

“Combination of security groups - rules for authorization of general ledger components” on page 154

In the Security Groups application, you combine the independent security groups or the security groups that are not independent to generate a security profile.

Authorizing standard services for security groups

When you assign a signature option (SIGOPTION) to a standard service, only users or groups that are authorized for that option can execute the standard service.

About this task

Standard service authorization does not support the use of conditions that may be associated with the signature option. Any condition assigned will be ignored.

Procedure

1. From the **Select Action** menu in the Security Groups application, choose **Standard Service Authorization**.
2. In the Standard Service Authorization window, click **New Row**.
3. Specify values for service, method, and option. After you specify these values, the **Application** field is populated with a value.
4. Click **OK**.

Results

The standard service is assigned a signature option. Any user that calls the standard service by means of the integration framework (HTTP, EJB, or SOAP) requires their authorization to the signature option to execute the standard service.

Related concepts:

“Combination of security groups” on page 152

In the Security Groups application, you can combine security groups to manage the security infrastructure within or across organizations.

“Combination of security groups - rules for application authorization” on page 153

In the Security Groups application, you combine the independent security groups or the security groups that are not independent to generate a security profile for a user.

“Combination of security groups - rules for site authorization” on page 155

In the Security Groups application, you combine the independent security groups or the security groups that are not independent to generate a security profile.

Overriding password duration for security groups

You specify password duration at the system level using the **Security Controls** action. However, you can make password duration specific to a security group by overriding the system setting.

About this task

If your implementation uses an application server to authenticate with an external directory (by means of LDAP), you cannot use this function because the directory manages passwords.

Procedure

1. In the Security Groups application, select the group for which you want to override the password duration.
2. From the **Select Action** menu, select **Override Password Duration**.
3. Specify the number of days that the password is to be valid for this group.
4. Specify the number of days before users in the group are alerted that the password is set to expire.
5. Click **OK**.

Related concepts:

“Combination of security groups” on page 152

In the Security Groups application, you can combine security groups to manage the security infrastructure within or across organizations.

“Combination of security groups - rules for application authorization” on page 153

In the Security Groups application, you combine the independent security groups or the security groups that are not independent to generate a security profile for a user.

“Combination of security groups - rules for site authorization” on page 155

In the Security Groups application, you combine the independent security groups or the security groups that are not independent to generate a security profile.

Specifying restrictions for security groups

You can specify all types of restrictions for security groups, such as data restrictions and collection restrictions.

Specifying data restrictions for security groups

To customize your security settings, you can specify restrictions on which records a security group can access. You can use a conditional expression or a conditional class file to define and apply these restrictions.

Procedure

1. In the Security Groups application, select the group for which you want to set restrictions.
2. On the **Data Restrictions** tab, select the type of restriction:
 - To specify restrictions on objects, click **Object Restrictions**.
 - To specify restrictions on attributes, click **Attribute Restrictions**.
3. Click **New Row**.
4. In the **Object** field, specify the table or view on which to set the restriction.
5. Optional: If you are specifying an attributes restriction, specify the attribute that you want to restrict.
6. Optional: In the **Application** field, specify the application to which you are applying the restriction. Leave the field blank to apply the restriction to all applications that use the object or attribute.
7. Specify the type of restriction.
8. Optional: Specify the following options for restrictions:

Option	Description
Reevaluate	Select this option for the restriction condition to be reevaluated when the user tabs to another field. If you do not select this option, the restriction conditions are evaluated after the changes to a field are saved.
Condition	Specify a conditional expression.

9. Save your changes.

Related concepts:

“Combination of security groups” on page 152

In the Security Groups application, you can combine security groups to manage the security infrastructure within or across organizations.

“Combination of security groups - rules for data restrictions” on page 152

When you combine independent or non-independent security groups, you use restrictions to specify the records that are visible to members of a security group.

Specifying collection restrictions for security groups

You can use data restrictions to specify which collections a security group can access.

Procedure

1. In the Security Groups application, select the group for which you want to set restrictions.
2. Click the **Data Restrictions** tab.
3. Click the **Collections Restrictions** tab.
4. Click **New Row**.
5. Specify the name of the collection to which you want to grant access.
6. Optional: Provide a description of the collection.
7. Save your changes.

Results

When you specify a collection restriction, a series of associated object restrictions is created. You can view this read-only series on the **Object** tab.

Related concepts:

“Combination of security groups” on page 152

In the Security Groups application, you can combine security groups to manage the security infrastructure within or across organizations.

“Combination of security groups - rules for application authorization” on page 153

In the Security Groups application, you combine the independent security groups or the security groups that are not independent to generate a security profile for a user.

“Combination of security groups - rules for site authorization” on page 155

In the Security Groups application, you combine the independent security groups or the security groups that are not independent to generate a security profile.

Specifying global data restrictions for security groups

Using global data restrictions, you can use a condition to specify restrictions on which records can be accessed by security groups.

Procedure

1. Optional: In the Conditional Expression Manager application, create one or more conditions to be evaluated to control access.
2. From the **Select Action** menu in the Security Groups application, select **Global Data Restrictions**.
3. In the Global Data Restrictions window, select one of the following options:
 - Specify an object restriction.
 - Specify an attribute restriction.
4. Click **New Row**.
5. Specify a value for the object or attribute. The object is the table or the view on which you set the restriction.
6. In the **Type** field, select the type of restriction that you want to set:
 - Select the application you want to apply the restriction to.
 - Leave the **Type** field blank for the restriction to apply to all applications that use the object or attribute.
7. Optional: If you want the restriction condition to be reevaluated when a user navigates away from a field, select the **Reevaluate** check box. If you do not select this check box, restriction conditions are still reevaluated, but only after you save the changes that you make to a field.
8. Optional: Specify a condition.
9. Click **OK**.

Related concepts:

“Combination of security groups” on page 152

In the Security Groups application, you can combine security groups to manage the security infrastructure within or across organizations.

“Combination of security groups - rules for data restrictions” on page 152

When you combine independent or non-independent security groups, you use restrictions to specify the records that are visible to members of a security group.

Specifying purchasing limits and tolerances for security groups

For security groups, you can specify approval limits for purchase requests, purchase orders, material requisitions, invoices, and contracts. You can also specify the amount that invoices, taxes, and services can deviate from an initial agreement.

Procedure

1. In the Security Groups application, select the group for which you want to set limits and tolerances.
2. Click the **Limits and Tolerances** tab.
3. Click **New Row**.
4. Specify an organization. After you specify an organization, the **Base Currency** field, which is read-only, is populated with the base currency 1 for each organization. You must use this currency in setting limits.
5. Optional: Specify the values for limits. When you create a record, the values in these fields default to 0. A value of 0 in a field indicates that the group has a limit of 0. A blank field means that the group has unlimited approval permissions.

Option	Description
PR Limit	Specify the maximum amount on a purchase request that the group can approve.
PO Limit	Specify the maximum amount on a purchase order that the group can approve.
MR Limit	Specify the maximum amount on a material requisition that the group can approve.
Invoice Limit	Specify the maximum amount on an invoice the that the group can approve.
Contract Limit	Specify the maximum amount on a contract that the group can approve.

6. Optional: Specify the upper and lower tolerances for invoices, taxes, and services. Use by amount or percent.
7. Save your changes.

Related concepts:

“Combination of security groups - rules for approval limits and tolerances” on page 153

In the Security Groups application, you combine the independent security groups or the security groups that are not independent to generate a security profile.

Deleting users from security groups

The security group to which a user belongs controls the level of access and privileges for the user within the system. To address changes in your security needs, you can remove a user from a security group.

Before you begin

You must delete a security group from the application server before you can remove users from a security group. To remove a user from a group, you must be authorized to reassign users to that group. Use the **Authorize Group Reassignment** action in the Security Groups application or the Users application to grant this authorization.

About this task

If you delete a security group in the directory server, but you do not want to delete the group from the system, you can delete the users from that group. The users will then no longer have access to the applications to which the group previously had access. If your implementation uses an application server to authenticate with a directory, you can associate users with groups in the directory or in the system, depending on your settings. You cannot delete users from the MAXEVERYONE group.

Procedure

1. In the directory server, delete the users from the security group that you want to delete. You must wait until the VMMSYNC cron task has fully synchronized the users and security groups in the directory server with the users and groups in the system.
2. Delete the security group in the directory server.
3. In the Security Groups application, select the group containing the relevant user.
4. Click the **Users** tab.
5. Delete the user.
6. Save your changes

Related concepts:

“Security groups and access to sites and applications” on page 140

Security access is based on security groups. You configure security groups to provide narrow access or broad access to applications, sites, and labor. You can also provide access to general ledger components, approval limits and tolerances.

Deleting security groups

As your business needs change, you can delete security groups. You can delete security groups in the Security Groups application, set up an archiving process, or create a cron task to remove the groups.

Before you begin

You must delete a security groups from the application server before you can delete the security group from the system.

About this task

You can delete a security group in the directory server but the VMMSYNC cron task does not delete the group from the system tables. If there are users associated with a security group or if a security group is specified as the default security group for all users, that security group cannot be deleted. You cannot delete the MAXEVERYONE group.

Procedure

1. In the directory server, delete the users from the security group that you want to delete. You must wait until the VMMSYNC cron task has fully synchronized the users and security groups in the directory server with the users and groups in the system.
2. Delete the security group in the directory server.
3. In the Security Groups application, select the group that you want to delete.
4. From the **Select Action** menu, select **Delete Group**.
5. Click **Yes**.

What to do next

If implementation uses an application server to authenticate with an external directory (by means of LDAP), a user that is deleted in the directory is not automatically deleted in the system. You must manually delete it.

Related concepts:

“Security groups and access to sites and applications” on page 140

Security access is based on security groups. You configure security groups to provide narrow access or broad access to applications, sites, and labor. You can also provide access to general ledger components, approval limits and tolerances.

Encrypting properties for security

You can encrypt properties in the `maximo.properties` file to provide additional security. The properties that are encrypted have `maxprop.encrypted=1` on the database.

Before you begin

The `maximo.properties` file is in the `<Maximo root> \applications\Maximo\properties` folder.

About this task

When you encrypt a property, the unencrypted original file remains on the file system as `maximo.properties_orig`. For security purposes, store the unencrypted original outside the system file structure.

Procedure

1. Access `maximo.properties` in a text editor.
2. Open a command shell and go to `<Maximo root> \tools\maximo` folder.
3. Type `encryptproperties` to run the batch file. The old file is renamed with an `*_orig` extension; for example, `maximo.properties_orig`.
4. Confirm that the new file contains an encryption string at the end.
5. Store the unencrypted originals (with the `*_orig` extension) outside the system file structure.

Related concepts:

“Encryption and security” on page 150

The data types Crypto, and CryptoX are used to encrypt passwords and other types of confidential information. The Java Cryptography Extension (JCE) is used to perform encryption.

Changing encrypted files for security

For security purposes, you can edit a file that you already encrypted.

Procedure

1. Delete the encrypted `maximo.properties`.
2. Restore the unencrypted originals back into the `<Maximo root>\applications\Maximo\properties` folder.
3. Remove the `_orig` extensions from the file.
4. Make your changes, then re-encrypt the file.

Related concepts:

“Encryption and security” on page 150

The data types Crypto, and CryptoX are used to encrypt passwords and other types of confidential information. The Java Cryptography Extension (JCE) is used to perform encryption.

Configuring the system to use application server security

You must configure the system to use application server security. However, if application server security was configured through the installer, you do not have to complete this procedure.

Before you begin

You must create the following items:

- A user directory on a Lightweight Directory Access Protocol (LDAP) server
- An organizational unit for the system
- A group, `maximousers`, under the organizational unit
- `MAXADMIN`, `MAXREG`, and `MXINTADMIN` administrative users in the directory that are assigned to an organizational unit and to the `maximousers` group

Procedure

1. In the System Properties application, set the value of the `mxe.useAppServerSecurity` property to 1.
2. For each `web.xml` file that sets the `useAppServerSecurity` value, modify the XML code:
 - a. Uncomment one of the `<login-config>` sections for FORM or BASIC login.
 - b. Set the value of `<useAppServerSecurity>` to 1.
 - c. Uncomment the `<security-constraint>` section.
3. To build the EAR file, change the directory to your `install_home\deployment` folder, and specify `buildmaximoear`.
4. Deploy the EAR file in the appropriate application server.
5. Synchronize the users and groups from LDAP into the system using the cron task.

What to do next

To configure to use application server security, you must configure your application server for LDAP security.

Related concepts:

“Security roles for the application server” on page 156

When you enable application server security, you can use roles to manage security.

“Single sign-on environment for application server security” on page 157

When you enable application server security, you can use a single sign-on environment. This environment enables a user to provide one name and password to access multiple applications.

“LDAP data synchronization” on page 158

Data synchronization keeps system data current with data in the Lightweight Directory Access Protocol (LDAP) directory server. Synchronized data moves only from the LDAP directory server to the system.

“Application server security - properties for user and group management” on page 156

After you enable application server security, you can configure the properties for group and user management to define whether the directory owns group creation, or user creation and management.

Configuring WebSphere Application Server for LDAP security

To implement configuration to use application server security, you configure your application server, configure to use Lightweight Directory Access Protocol (LDAP).

.

Before you begin

You must configure your system for application server security.

Procedure

1. Complete the procedure to configure WebSphere Application Server to use LDAP security. For specific instructions on how to configure WebSphere Application Server, see the WebSphere Application Server Information Center, at  <http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/index.jsp>, and search for WebSphere Active Directory.
2. Restart the application server.
3. Deploy the system and map the security role, maximouser, to the users and groups that meet the requirements for your organization, or assign the users to the default group, maximousers, in the LDAP system.

Related concepts:

“Security roles for the application server” on page 156

When you enable application server security, you can use roles to manage security.

“Single sign-on environment for application server security” on page 157

When you enable application server security, you can use a single sign-on environment. This environment enables a user to provide one name and password to access multiple applications.

“LDAP data synchronization” on page 158

Data synchronization keeps system data current with data in the Lightweight Directory Access Protocol (LDAP) directory server. Synchronized data moves only from the LDAP directory server to the system.

“Application server security - properties for user and group management” on page 156

After you enable application server security, you can configure the properties for group and user management to define whether the directory owns group creation, or user creation and management.

Configuring two directory servers

Two directory servers can be configured for the deployment.

About this task

You can configure the Virtual Member Manager to use two separate directory servers to authenticate users that log in to Maximo Asset Management. You might choose this approach if you already have more than one directory server in your environment, or if you do not want to include systems users such as wasadmin in your enterprise directory server. You can use any combination of supported directory servers: two instances of IBM Tivoli Directory Server, two instances of Microsoft Active Directory, or one of each.

The two directory servers must be defined in the same realm. There must not be any user name that appears in both directory servers. If, after performing this configuration, you create users using the WebSphere interface, they will be defined in the first directory server that you configure the Virtual Member Manager to use.

The following procedure assumes that both directory servers have been installed. If you choose to have Maximo Asset Management configure a directory server during the installation process, then you can consider it your first directory server and configure the Virtual Member Manager to add a second directory server. If you are reusing existing directory servers, or if you have installed a new directory server but did not configure it using Maximo Asset Management installation program, then you must complete all the steps to configure both servers.

To use two separate directory servers to authenticate Maximo Asset Management users, you must configure the Virtual Member Manager to federate both directory servers, and you must configure cron tasks to synchronize the Maximo user directory with both directory servers. To accomplish these tasks, follow these steps:

Procedure

1. Configure the Virtual Member Manager to use the first directory server. If you choose to have Maximo Asset Management configure a directory server during the installation process, skip this step.

2. Follow the same set of steps to configure the Virtual Member Manager to use the second directory server. Be sure that the Realm name value is the same as the value for the first directory server. Give the second directory server a different Repository identifier.
3. After the installation of Maximo Asset Management is complete, log in to the Maximo Asset Management interface, and click **Go To > System Configuration > Platform Configuration > Cron Task Setup**.
4. Type VMM in the **Cron Task** field, and press **Enter**.
5. Locate the VMMSYNC cron task, and click it.
6. Set the task to **active**. This task completes the configuration of the cron task for the first directory server.
7. Duplicate the existing VMMSYNC cron task and modify these fields:
 - Group Mapping
 - User Mappings

Ensure that the BaseDN for both group mappings and user mappings parameters is provided correctly. The BaseDN value instructs the VMMSync crontask to search for users and groups in a particular location in the directory server. The BaseDN value for user and group always ends with the base entry value provided when federating the directory server under Virtual Member Manager.

The principal and credential values must remain the same as the first directory server.
8. Schedule the task to run every 5 minutes (or a different interval if you prefer), set it to active, and set it to keep history records.
9. Save the task.
10. After 5 minutes, check whether the task has run and restart it if it has not.
11. After both cron tasks have run, click **Go To > Security > Users** and verify that users from both directory servers appear in the list.

Configuring WebLogic Server for LDAP security

To implement configuration to use application server security, you configure your application server to use Lightweight Directory Access Protocol (LDAP).

Before you begin

You must configure your system for application server security.

Procedure

1. Complete the procedure to configure WebLogic Server to use LDAP security. For specific instructions on how to configure WebLogic Server, see the WebLogic Server documentation, and search for WebLogic Active Directory.
2. Restart the application server.
3. Deploy the system and map the security role, maximouser, to the users and groups that meet the requirements for your organization, or assign the users to the default group, maximousers, in the LDAP system.

Related concepts:

“Security roles for the application server” on page 156

When you enable application server security, you can use roles to manage security.

“Single sign-on environment for application server security” on page 157

When you enable application server security, you can use a single sign-on environment. This environment enables a user to provide one name and password to access multiple applications.

“LDAP data synchronization” on page 158

Data synchronization keeps system data current with data in the Lightweight Directory Access Protocol (LDAP) directory server. Synchronized data moves only from the LDAP directory server to the system.

“Application server security - properties for user and group management” on page 156

After you enable application server security, you can configure the properties for group and user management to define whether the directory owns group creation, or user creation and management.

Changing cron task parameters for data synchronization

When you synchronize data from the Lightweight Directory Access Protocol (LDAP) directory server to the system, the application server synchronizes on the common name (cn) attribute. The attribute is the **Full Name** field in the LDAP directory server. You can change the default values of the parameters for synchronization if necessary.

About this task

If you want to synchronize from the user logon name and log on to the system, the user name attribute must be correctly mapped. You must map the attribute in the LDAPSYNC cron task parameters, the VMMSYNC cron task parameters, and in the application server. The field length in the directory must be the same value as the field length in the system tables. If this value is smaller in the system tables, you can increase the maximum length of the field in the Database Configuration application.

Procedure

1. Open the Cron Task Setup application and select the appropriate cron task:

Option	Description
If you use Active Directory and this directory is the only directory you use regardless of the application server you use	Select the LDAPSYNC cron task
If you use Active Directory, a directory platform or directory instance, and another directory, only Virtual Member Manager (VMM) and IBM WebSphere Application Server support these directories.	Select the VMMSYNC cron task

2. On the **Parameters** tab of the **Cron Task** tab, select the UserMapping parameter and review the value.
3. If the value of the UserMapping parameter is smaller in the system tables, select the LDAPSYNCCRONPARM object in the Database Configuration application.
4. On the **Attributes** tab, select the UserMapping attribute, increase the length of the field, and save the attribute.

5. Repeat steps 5 and 6 for the VMMSYNCCRONPARM object.

What to do next

You must ensure the cron tasks are active and then synchronize the data.

Related concepts:

“LDAP data synchronization” on page 158

Data synchronization keeps system data current with data in the Lightweight Directory Access Protocol (LDAP) directory server. Synchronized data moves only from the LDAP directory server to the system.

Activating cron tasks to synchronize data

Synchronization keeps data in the system current with data in the Lightweight Directory Access Protocol (LDAP) directory server. You activate cron tasks to synchronize the data.

Before you begin

You must log on to the system as an administrative user. The only directories that are supported are IBM Tivoli Directory Server and Microsoft Active Directory.

Procedure

1. Open the Cron Task Setup application and select the appropriate cron task:

Option	Description
If you use Active Directory and this directory is the only directory you use regardless of the application server you use	Select the LDAPSYNC cron task
If you use Active Directory, a directory platform or directory instance, and another directory, only Virtual Member Manager (VMM) and IBM WebSphere Application Server support these directories.	Select the VMMSYNC cron task

2. In the Cron Task Instances section of the **Cron Task** tab, select the **Active** check box.
3. Specify a schedule for the cron task and save the changes. Allow the cron task run to synchronize all the users and groups from the LDAP directory server into the database tables.

Related concepts:

“LDAP data synchronization” on page 158

Data synchronization keeps system data current with data in the Lightweight Directory Access Protocol (LDAP) directory server. Synchronized data moves only from the LDAP directory server to the system.

Configuring WebSphere Application Server for incremental synchronization

You can configure WebSphere Application Server to use incremental synchronization in addition to the full synchronization mode. Incremental synchronization is a more efficient method of updating user data than full synchronization because as only data that has changed is synchronized from the directory to the tables.

Before you begin

You must configure WebSphere Application Server for incremental synchronization before you configure Maximo Asset Management for incremental synchronization. In the WebSphere Administration Server wsadmin tool, set the **supportChangeLog** parameter to native. The underlying repository that you use must support change polling.

About this task

For information about the wsadmin tool, see the IBM

WebSphere Application Server documentation and search for WebSphere Application Server administration.

Procedure

1. Open the Cron Task Setup application and select the VMMSYNC cron task.
2. Select the **ChangePolling** parameter and specify the value as 1.
3. Save the cron task.
4. From the **Select Action** menu, click **Reload Request**.

Related concepts:

“LDAP data synchronization” on page 158

Data synchronization keeps system data current with data in the Lightweight Directory Access Protocol (LDAP) directory server. Synchronized data moves only from the LDAP directory server to the system.

Setting password requirements

You can set password requirements, set password characters, allow the placement of password characters, and create and delete excluded passwords. When the system is implemented, the initial password requirement is of a minimum length. You can change the minimum length and specify additional requirements.

About this task

Your implementation might use an application server to authenticate with a directory. In that case, you do not use the system to perform some functions. These functions are performed in the directory and synchronized into the system. These functions include:

- Adding users (including self-registration)
- Adding security groups
- Associating users with security groups
- Managing passwords

Procedure

1. From the **Select Action** menu, select **Security Controls**. You use the Security Controls window to perform the following functions:

Option	Description
To set password requirements	1. In the Password Requirements section, in the Minimum Password Length field, set the minimum length for a user password. The default minimum password length is 6 characters. 2. In the Number of Identical Adjacent Characters Allowed in Password field, set the number of identical adjacent characters that are allowed in a password. 3. Select the Password can Contain Login ID check box to allow a user to use a login ID in a password.
To set required password characters	1. In the Required Password Characters section, set the following specifications: • Must Include an Uppercase Character - If you select this check box, the password must contain at least one uppercase character. • Must Include a Lowercase Character - If you select this check box, the password must contain at least one lowercase character. • Must Include a Number - If you select this check box, the password must contain at least one numeric character. Numeric characters are: 1 2 3 4 5 6 7 8 9 0. • Must Include a Special Character (!, @, #, etc.) - If you select the check box, the password must contain at least one special character. Supported special characters are: ! @ # \$ % ^ & * () - _ = + \ [] { } ; : / ? . > <

Option	Description
To allow placement of password characters	- In the Allowed Placement of Password Characters field, set the following specifications: First Character can be a Number - If you select this check box, the first character of a password can be a number. Last Character can be a Number - If you select this check box, the last character of a password can be a number. Numeric characters are: 1 2 3 4 5 6 7 8 9 0. First Character can be a Special Character - If you select this check box, the first character of a password can be a special character. Supported special characters are: ! @ # \$ % ^ & * () - _ = + \ [] { } ; : / ? . > < . Last Character can be a Special Character - If you select this check box, the last character of a password can be a special character. Supported special characters are: ! @ # \$ % ^ & * () - _ = + \ [] { } ; : / ? . > < .
To create an excluded password list	- In the Excluded Password List section, click New Row. - In the Password field, type the password that you want to prohibit from being used on the system.
To delete an excluded password	In the Excluded Password List section, click Mark for Delete next to the excluded password that you want to delete.

2. Click **OK**.

Generating passwords

You can configure the asset management system to generate user passwords. These generated passwords can then be sent as e-mail notifications to asset management system users.

About this task

If your implementation uses an application server to authenticate an external directory (LDAP), do not use the asset management system to set or change passwords and password hints. All password-related functions are managed by the directory.

Procedure

- From the **Select Action** menu in the Users or the Security Groups application, select **Security Controls**.
- In the **Template for Emailing Reset Passwords** field, enter a template value.
- In the Automatic Password Generation table window, select one of the following options:

- Always Email Generated Passwords to Users (Never Display Screen)
- Allow Generated Passwords to Be Displayed On Screen

4. Click OK.

Enabling login tracking

You use the **Security Controls** action in the Security Groups and Users applications to enable login tracking. Login tracking enhances security by limiting the number of incorrect passwords a user can enter when attempting to sign in.

Before you begin

You must enable login tracking to control the number of login attempts allowed for a user. If you set the number of login attempts without enabling login tracking, the asset management system does block the user when the number of maximum number of login attempts is exceeded.

About this task

When you enable login tracking, the asset management system logs all sign in attempts, successful and unsuccessful. You can specify the maximum number of unsuccessful sign in attempts. If a user exceeds the maximum number of unsuccessful attempts, the status of the user record is changed to **BLOCKED**. The user is prevented from logging in until an administrator uses the **Change Status** action to set the status back to **ACTIVE**.

Procedure

1. From the Select Action menu in the Security Groups or Users application, select **Security Controls**.
2. In the Security Controls window, specify whether you want to enable the login tracking:

Option	Enable Login Tracking
Enabled login tracking	Selected
Disabled login tracking	Cleared

3. Enter values in the following fields:

Option	Description
Login Attempts Allowed	The number of times a user can incorrectly enter their user name or password before being blocked from the asset management system.
Password lasts this Number of Days	The number of days a user password is valid. If you do not want the password to expire, leave this field empty.
Days Before Password Expires to Warn User	The number of days before the user receives an expiration warning message.
Days Before Previously Used Password can be Reused	The number of days before a user can use a previously used password. If this value is 0, the asset management system does not check for password reuse.

4. Click OK.

Chapter 7. Registering users

User records contain user names, passwords, and security profiles that determine the applications, options, and data to which a user can access.

About this task

To manage security privileges, you can perform the following functions for users:

- Manage user status.
- View user security profile.
- Specify various user defaults, such as default insert site, default storeroom, default language, and default general ledger accounts for purchasing. The default purchasing account is the general ledger account that is used for desktop requisitions, but not for all purchasing.
- Grant users the right to access inactive sites.
- Specify which users can access a screen reader to assist in interacting with the system.
- Set system-wide security controls and new user default groups.
- Change passwords (if you are using an external directory, this functionality is not available).
- Create database users (if you are using an external directory, this functionality is not available).

Users overview

You can use the Users application to manage users. User records contain user names, passwords, and security profiles that determine the applications, options, and data to which a user can access.

Administrative users

Administrative users have full or restricted access to the Security Groups application and to the Users application. Administrative users are responsible for implementing and maintaining security services, such as adding users, building profiles, or managing general site administration.

Administrative users might need access to the following applications to perform system administration"

- Actions
- Application Designer
- Calendars
- Chart of Accounts
- Classifications
- Communication Templates
- Cron Task Setup
- Database Configuration
- Currency Codes
- Domains

- E-mail Listeners
- Escalations
- Exchange Rates
- Integration
- Launch in Context
- Logging
- Object Structures
- Organizations
- Roles
- Security Groups
- Sets
- System Properties
- Users
- Web Services Library
- Workflow Administration
- Workflow Designer

Some users might assign management functions that are administrative in nature to supervisors or managers, especially in the areas of Information Technology Asset Management and Service Desk operations. These users are not considered administrative users.

Related tasks:

“Adding users” on page 193

To manage users, you can create records that contain user names, passwords, and security profiles. These records determine which applications, options, and data that user can access.

“Assigning users to security groups” on page 194

To manage security settings and to grant user privileges, you can assign users to security groups. New users are assigned to the default group (MAXDEFLTREG) and the group for all users (MAXEVERYONE). The default group is used to give newly registered users basic privileges, and the group for all users is used to specify global settings.

“Changing persons associated with users” on page 195

You can manage your work force information by associating user IDs with specific person records. Once an association is no longer being used, you can reuse the user ID with another person record.

“Changing user settings” on page 197

You can change user settings, such as the settings for storerooms and insert sites, and settings for screen reader access.

Database users

Database users are granted access to read, insert, update, and delete specific objects that define a set of fields and business rules. These objects can also update one or more database tables.

In the Users application, you use the **Database Access** action to create database users. If you implement databases that need operating system IDs, you must create the operating system ID.

The default database username is maximo.

The following commands detail the standard authorizations that are required:

- Create user *dbusername* identified by the system
- Alter user *dbusername* default table space quota as unlimited
- Alter user *dbusername* temporary table space temp
- Grant create trigger to *dbusername*
- Grant create session to *dbusername*
- Grant create sequence to *dbusername*
- Grant create synonym to *dbusername*
- Grant create table to *dbusername*
- Grant create view to *dbusername*
- Grant create procedure to *dbusername*
- Grant alter session to *dbusername*
- Grant execute on ctxsys.ctx_ddl to *dbusername*

To allow database access to users, the following commands detail the additional grants that are required to create database users:

- Grant create user to *dbusername*
- Grant drop user to *dbusername*
- Grant create session to *dbusername* with admin option
- Grant alter user to *dbusername*

System users

There are certain user IDs, like MAXADMIN and MAXREG, that are required for the system to run properly. These user IDs are known as system users

The system users MAXADMIN and MAXREG are part of the database. To create system users, you use the system account option in the Users application. You cannot delete a system user.

Configuration of self-registration for users

To allow users to self-register, you can configure the self-registration process.

The following list describes configurations for the self-registration process:

- If wanted, in either the Users application or Security Groups application, use the **Security Controls** action to rename the default group for new users and the initial self registered user status. The default values for these items are MAXDEFLTREG and NEWREG, respectively.
- If you want self-registered users to have basic authorizations, in the Security Groups application, add the wanted authorizations to the default group. By default, the authorizations for this group are limited to changing an expired password, and accessing the start center.
- Self-registration requests are placed in the SELFREG workflow process. By default, this process notifies an administrator when a new self-registration is pending. The administrator must approve the new user before the user can access the system. You can change this process to be specific to your organization. For example, you can change this process to automatically approve self-registered users.
- When a self-registration request is being processed, email notifications are sent as the self-registration progresses. These notifications are available as communication templates. You can change the content of these templates to be

specific to your organization. The templates include: NEWSSELFREG, REGNOTIFY, REGAPPROVE, and SELFREGREJ.

- In the Users application or in the Security Groups application, you can use the **Security Controls** action to specify the following statuses for self-registered users:
 - NEWREG - This value is the default status for self-registered users. Users with this status cannot log in, and workflow is enabled.
 - ACTIVE - Users with this status can log in, and workflow is disabled.
 - INACTIVE - Users cannot log in, and workflow is disabled.
- In the Users application, you can use the workflow process to process self-registered users. You enable the workflow process in the Security Groups application. To configure and activate the workflow process, use the Workflow Designer application.

Self-registration for users

New users can use the self-registration process to register with a minimum amount of information.

To self-register, use the **register now** function on the home page. Self-registered users are assigned to a default security group, MAXDEFLTREG. In the Users application or Security Groups application, an administrative user can replace the default security group with another security group. For example, you can configure the access rights and privileges of the default group for self-registered users to reflect the business rules for your company.

Additional settings are made after self-registration requests are routed to an administrator through workflow (if enabled). The requests are either approved or rejected. After being approved, the administrator assigns self-registered users to appropriate security groups. The administrator also notifies the users that they can use the system. The user must provide additional information using the Profile link in the navigation bar.

After submitting a self-registration request, a person record and user record is created for the user.

The following table shows the required and optional information a user provides to create a self-registration request.

Table 35. User self-registration information and optional information

User self-registration information	Optional information
First name	Password hint question
Last name	Answer
User name	Supervisor
Password	Default insert site
Confirm password	Default storeroom
Primary email	Primary phone
	Language
	Locale
	Time zone
	Additional information

The following table lists the registration information that is hidden from the user.

Table 36. User self-registration hidden information

Hidden information	Setting
User status	Defaults based on the REGSTATUS setting in the MAXVAR table
Person ID	Defaults to the user ID
User name	Defaults to the user ID
Force password expiration	Defaults to yes for new users that are created using the Users application or self-registration, and when an administrator changes a password for a user
Query with site	Defaults to yes
Person status	Defaults to active
Transaction notifications	Defaults to never
Workflow notifications	Defaults to process
Accepting workflow email	Defaults to yes

Security controls

You can implement security controls, such as setting passwords and password hints for users, and creating security profiles.

Related tasks:

“Implementing security for users” on page 199

You must implement security for users, such as specifying security groups, specifying passwords, and specifying security profiles.

Passwords for users

For security purposes, users must define a password before logging in.

As an administrator, you can change passwords for users. You can use the password hint question and answer to verify that the person requesting a password is the correct user. As a user, you can set password hints using the **Set Password Hint** action. Authorized users also can change their own password and set their password hint question and answer using the **Password Information** action in the **Profile** link.

When you use an application server for authentication, the directory manages user passwords and all password-related functions.

Password expiration

In the Users application, you can use the **Security Controls** action to specify settings for password expiration. These settings include the number of days for the duration of a password, warning of password expiration, and the number of days before a password can be reused.

Automatic passwords

In the Users application, you can use the **Security Controls** action to specify the random generation of passwords. After the password is generated, an email notification is sent to the appropriate user. You can also specify that generated

passwords are either sent in an email notification to the user or that generated passwords are displayed on the screen.

Password requirements

In the Users application, you can use the **Security Controls** action to specify the following password requirements:

- Minimum password length.
- The number of identical adjacent characters.
- If a password can contain a login ID.
- Required password characters, such as uppercase and lowercase characters, numbers, and special characters.
- Allowed placement of password characters, such as whether the first or last character can be a number, or if the first or last character can be a special character.

Excluded password list

In the Users application, you can use the **Security Controls** action to manage a list of excluded passwords.

Database passwords (Oracle and SQL Server only)

If the user is authorized to access the database, the Database Password section displays their database user ID. You can use **Set Password** to change the database password.

You can select the **Also Change Database Password to This Password** check box to change the database password to match their password. If you change a database password to match a system password, the password must satisfy the requirements of both passwords. For example, if you use the Oracle database platform, the password must meet the special character requirements that Oracle supports.

Related tasks:

“Changing system and database passwords for users” on page 199

For security purposes, you can change a system password and also change their database password to synchronize with the new system password. Authorized users can change their own passwords.

“Specifying passwords for new users” on page 199

If you provide a user name and email address, you can specify or generate the initial system password for a new user.

“Specifying password hints for users” on page 200

You can specify password hints to check if the person requesting a new password is the user.

Password hints for users

Password hints help users remember the passwords that they created and they require users to confirm their identity before an administrator resets passwords. Password hints also help users when they reset passwords.

Remembering passwords

Users can use the **Password Information** action, available from the **Profile** menu, to set their own password, password hint question, and password hint answer.

When users reset their password, they can select a password hint and provide the answer.

Confirming identity

Administrators and help desk agents can require users to answer the password hint question correctly before they can reset password. Administrators and help desk agents can also change the password hint question and answer for a user.

Resetting passwords

Before logging in, users can use the **forgot your password** link to reset passwords by providing their password hint question, email address, and password hint answer.

Security authorizations for users

You can define for users the security authorizations to applications, tabs, actions, and fields. After you create security groups and define their security authorizations, you create user records and assign each user to one or more security groups.

By default, when you use an application server for authentication, the directory manages user creation. You can set properties to let user creation be performed directly in the system. The settings of these properties result in certain features being enabled or disabled in the system.

Related tasks:

“Implementing security for users” on page 199

You must implement security for users, such as specifying security groups, specifying passwords, and specifying security profiles.

Security profiles for users

You can specify the security profile for a user. You can assign the user to groups and set other security attributes. These attributes include default insert sites, storeroom sites, and default storerooms.

The following points are applicable to security profiles:

- The default insert site is the default value to use when a user creates a record. The value in the **Organization** field is the organization to which the default insert site belongs.
- By default, new users are assigned to the MAXDEFLTREG group. You cannot delete this group from a security profile.
- If the use default insert site as a display filter option is selected, the user views records for only the default insert site on the **List** tab of all applications.
- The value for default storeroom applies to when a user creates a material requisition. You can specify a value for the storeroom site only. However, if you specify a default storeroom, you must also specify the storeroom site. If you enter a default insert site or storeroom site, you must grant the user privileges to a group with that site.

Database access for users

By default, users do not have authorization to access the database. However, in some instances, users need access to the database. For example, a user might need database authorization to view tables and columns to create reports.

You can grant a user authorization to read, insert, update, and delete specific objects that define a set of fields and business rules, and that update one or more database tables. Before you can grant database access, you must complete additional steps to grant authorization to database users. To provide a user complete database authorization, use the tools and procedures of your database platform.

Database passwords

The database user ID and password do not have to be the same. However, the database password must meet the requirements of your database platform. For example, if you are using the Oracle database platform, you cannot create a password with certain special characters that Oracle does not support.

You use the **Change Password** action to change a password for database users. You can set a database user password to match a system password. However, the database user password must support the password requirements. Users cannot change their own database passwords.

Default insert sites for users

In the Users application, you assign users to a default insert site to insert records. The records that users view pertain only to the default insert site. Without a default insert site, some applications cannot function. For example, without a default insert site, users cannot add purchase orders in a site-level application.

For site-level records, the site defaults to the value of the default insert site. For organization-level records, the organization defaults to the organization of the site that is specified as the default insert site. The Profile link in the navigation bar enables users to change their default insert site.

The **Use Default Insert Site as a Display Filter** setting is a filter that allows users to view only records from their default insert site. However, if the application is at the organizational level, this setting allows users to view all records from all sites within the organization. If you do not use this setting, users can view records for all the sites to which they have access.

User statuses

User statuses determine how users can interact with the system.

A user record can have one of the following statuses:

Active The default status for new records. To log in, a user record must be active.

Blocked

Users cannot log in. When login tracking is enabled, a user can be blocked when they incorrectly enter their name or password beyond the number of times specified in security controls. Blocked is a system-generated status; you cannot use the Change Status action to change a status to blocked.

Deleted

The user was deleted but the user name has been retained because login tracking is enabled. You can configure login tracking in the Security Controls window. If login tracking is not enabled, all evidence of a user that is deleted is removed from the database and the user name can be reused. If login tracking is enabled, the user name is retained in the

MAXUSER table with a status of deleted and the user name cannot be reused. User names that have a status of deleted are not displayed in the Users application.

Inactive

Users cannot log in. Inactive user records do not appear in select value lists and cannot be associated with new records.

Newreg

The default status for user records that are created using self-registration. This status is used to identify user records to route into a workflow process.

When you use a directory server for authentication, the blocked status and newreg status are not available.

Working with users

You manage users by adding them, by assigning them to security groups, by changing general ledger accounts for users, and by performing other administrative tasks.

Adding users

To manage users, you can create records that contain user names, passwords, and security profiles. These records determine which applications, options, and data that user can access.

Before you begin

A user record must have a corresponding person record. You create person records either in the People application or in the Personal section in the Users application.

About this task

By default, when you use an application server for authentication, the directory manages user creation. You can set properties to let user creation be performed directly in the system. The settings of these properties result in certain features being enabled or disabled.

Procedure

1. In the User application, click **New User**.
2. Optional: In the **User** field, type a unique user identifier. If you use autonumbering, the **User** field is already populated.
3. Optional: If no matching person record is found, select one of the following options to create a person record:
 - To create a person record based on information in the Personal section of the user record, click **Yes**.
 - To select an existing person record, click **No**. In the **Person** field, select the person record.
 - To create a person record, click **No**. In the **Person** field, select **Detail Menu** and then select **Go to People**.
4. Optional: In the **User Name** field, type a value. The user name defaults to the ID name in the **User** field. This value is used to log in. You can change the user name to an employee number or email address.

5. Click **Set Password** to define a password.
6. Optional: Specify additional personal, user setting and general ledger account information.

Option	Description
Default Insert Site	As a best practice, assign a default insert site. Without a default insert site, many applications cannot function. When you change the default storeroom, you change the site from where the material requisitions for a user are fulfilled. When you change the default insert site, you assign a user a default insert site to insert records. The records that the user can view are only to the default insert site.
Language, Locale, Timezone	Specify details for the user.
System Account	Select the check box to create system users, such as MAXADMIN and MAXREG, which are required for the system to run properly
Can Access Inactive Sites	Select the check box to use inactive site access to set up or decommission sites. By default, the check box is cleared.
Use Screen Reader	Select the check box to allow a user to access a screen reader for accessibility reasons.
Organization and GL Account	Specify details to track expenses. When you change the default general ledger account, you change which account is charged or credited for financial transactions related to a specific record.

7. Save the record.

Results

The user is assigned to the default security group, MAXDEFLTREG.

What to do next

To have security authorizations, a new user must be assigned to security groups.

Related concepts:

“Administrative users” on page 185

Administrative users have full or restricted access to the Security Groups application and to the Users application. Administrative users are responsible for implementing and maintaining security services, such as adding users, building profiles, or managing general site administration.

Assigning users to security groups

To manage security settings and to grant user privileges, you can assign users to security groups. New users are assigned to the default group (MAXDEFLTREG) and the group for all users (MAXEVERYONE). The default group is used to give newly registered users basic privileges, and the group for all users is used to specify global settings.

About this task

By default, when you use an application server for authentication, the directory manages user creation. You can set properties to allow user profile be created in the system. You can add only security groups that you are authorized to manage. You can delete a security group from a user profile when the user is assigned to another security group.

Procedure

1. In the Users application, select the user to whom you would like to assign security groups and select the **Groups** tab.
2. In the Groups section, complete one of the following steps.

Option	Description
Click New Row	Specify a security group for a user
Click Select Groups	Specify multiple security groups

3. Click OK.

Related concepts:

“Administrative users” on page 185

Administrative users have full or restricted access to the Security Groups application and to the Users application. Administrative users are responsible for implementing and maintaining security services, such as adding users, building profiles, or managing general site administration.

Authorizing users to assign other users to security groups

When you create a security group, you are authorized to assign users to that group. As an administrator, you can authorize a user to assign other users to security groups.

Procedure

1. In the Users application, access the user whom you want to authorize.
2. From the **Select Action** menu, select **Authorize Group Reassignment**.
3. To authorize assignments to security groups, select one of the following options:
 - To authorize one security group, in the **Group** field, specify a group.
 - To authorize several security groups, click **Select Groups** and select the check boxes next to the groups that you want to add.
4. Click OK.

Changing persons associated with users

You can manage your work force information by associating user IDs with specific person records. Once an association is no longer being used, you can reuse the user ID with another person record.

Procedure

1. In the Users application, access the user for whom you want to make an association.
2. From the **Select Action** menu, select **Change Person**.
3. In the **New Person** field in the Change Person window, specify a new person record.

4. Click **OK**.

Example

For example, there is a user ID with the value of contractor 1 that is associated with Bob Smith from XYZ Consulting. The contract for Bob Smith expired. Therefore, you can associate the user ID, contractor 1, with another user ID.

Related concepts:

“Administrative users” on page 185

Administrative users have full or restricted access to the Security Groups application and to the Users application. Administrative users are responsible for implementing and maintaining security services, such as adding users, building profiles, or managing general site administration.

Changing the status of multiple users

You can either allow or prevent a user from logging in by changing the status for the user.

About this task

The default status of a user is active.

Procedure

1. From the **List** tab in the Users application, select the users whose statuses you want to change.
2. From the **Select Action** menu, select **Change Status**.
3. Select the new status.
4. Optional: In the **Memo** field, type the reason for the status change.
5. Click **OK**.

Example

When users exceed the number of allowed unsuccessful login attempts, their status is changed to inactive. To allow the users to log in, you change their status to active.

Related concepts:

“User statuses” on page 192

User statuses determine how users can interact with the system.

Changing the status of users

You can either allow or prevent a user from logging in by changing the status for the user.

About this task

The default status of a user is active.

Procedure

1. In the Users application, open the user whose status you want to change.
2. From the **Select Action** menu, select **Change Status**.
3. Select a new status.
4. Optional: In the **Memo** field, enter the reason for the status change.

5. Click **OK**.

Example

When a user exceeds the number of allowed unsuccessful login attempts, their status is changed to inactive. To allow the user to log in, you change the status to active.

Related concepts:

"User statuses" on page 192

User statuses determine how users can interact with the system.

Changing user settings

You can change user settings, such as the settings for storerooms and insert sites, and settings for screen reader access.

Related concepts:

"Administrative users" on page 185

Administrative users have full or restricted access to the Security Groups application and to the Users application. Administrative users are responsible for implementing and maintaining security services, such as adding users, building profiles, or managing general site administration.

Changing user settings for inactive site access

You use inactive site access to set up or decommission sites. By default, the setting to access an inactive site is off.

Procedure

1. In the Users application, access the user for whom you want to change the inactive site access setting.
2. In the User Settings section, select the **Can Access Inactive Sites** check box.
3. Save your changes.

Changing user settings for language, locale, and time zone

The language, locale, and time zone settings determine the display settings for users. You can change these settings to suit your needs.

Procedure

1. In the Users application, access a user record.
2. Change the value in the **Language** field, **Locale** field, or **Time Zone** field.
3. Save your changes.

Changing user settings for screen readers

For accessibility reasons, you can turn on the setting for a user to access a screen reader.

Procedure

1. In the Users application, open a user record.
2. In the User Settings section, select the **Use Screen Reader** check box.
3. Save your changes.

Changing user settings for storerooms and insert sites

When you change the default storeroom, you change the site from where the material requisitions for a user are fulfilled. When you change the default insert

site, you assign a user a default insert site to insert records. The records that the user can view pertain only to the default insert site.

Procedure

1. In the Users application, open the record for the user whose settings you want to change.
2. Change the values for one or more of these settings:
 - In the **Default Insert Site** field, specify a value.
 - In the **Storeroom Site for Self-Service Requisitions** field, specify a value.
 - In the **Default Storeroom for Self-Service Requisitions** field, specify a value.
3. Save your changes.

Related concepts:

“Default insert sites for users” on page 192

In the Users application, you assign users to a default insert site to insert records. The records that users view pertain only to the default insert site. Without a default insert site, some applications cannot function. For example, without a default insert site, users cannot add purchase orders in a site-level application.

Changing settings for storerooms and insert sites for multiple users

When you change the default storeroom, you change the site from where the material requisitions for a user are fulfilled. When you change the default insert site, you assign a user a default insert site to insert records. The records that the user can view pertain only to the default insert site.

Procedure

1. In the Users application, select the users whose settings you want to change.
2. From the **Select Action** menu, select **Set Security Profile**.
3. To change the settings, complete these steps:
 - a. In the Update User Defaults section, select the **Edit** check box for the setting that you are changing.
 - b. Change the value for the default insert site, select the default insert site as a display filter, change the default storeroom site for self-service requisitions or change the default storeroom site for self-service requisitions.
4. Click **OK**.
5. Save your changes.

Related concepts:

“Default insert sites for users” on page 192

In the Users application, you assign users to a default insert site to insert records. The records that users view pertain only to the default insert site. Without a default insert site, some applications cannot function. For example, without a default insert site, users cannot add purchase orders in a site-level application.

Changing general ledger accounts for users

General ledger accounts are used to track expenses. When you change the default general ledger account, you change which account is charged or credited for financial transactions related to a specific record.

Procedure

1. In the Users application, open the record for the user for whom you want to change the general ledger account.

2. In the **Organization** field in the Purchasing section, specify a value.
3. In the **GL Account** field, specify a value.
4. Save your changes.

Implementing security for users

You must implement security for users, such as specifying security groups, specifying passwords, and specifying security profiles.

Related concepts:

“Security controls” on page 189

You can implement security controls, such as setting passwords and password hints for users, and creating security profiles.

“Security authorizations for users” on page 191

You can define for users the security authorizations to applications, tabs, actions, and fields. After you create security groups and define their security authorizations, you create user records and assign each user to one or more security groups.

Specifying passwords for new users

If you provide a user name and email address, you can specify or generate the initial system password for a new user.

About this task

When you use an application server for authentication, the directory manages user passwords.

Procedure

1. In the User application, click **New User**.
2. In the New Password section in the Set Password window, generate a random password or specify a password.
3. To send an email message with the password to the user, select the **email Password to User** check box. This check box is read-only and is selected when the **Always email Generated Passwords to Users (Never Display on Screen)** check box in the Security Controls window is selected. This check box is editable when the **Allow Generated Passwords to Be Displayed on Screen** option in the Security Controls window is selected.
4. Optional: Specify that the user change their password when logging in for the first time.
5. Click **OK**.

Related concepts:

“Passwords for users” on page 189

For security purposes, users must define a password before logging in.

“Password hints for users” on page 190

Password hints help users remember the passwords that they created and they require users to confirm their identity before an administrator resets passwords. Password hints also help users when they reset passwords.

Changing system and database passwords for users

For security purposes, you can change a system password and also change their database password to synchronize with the new system password. Authorized users can change their own passwords.

Before you begin

When you use an application server for authentication, the directory manages user passwords.

About this task

Oracle and SQL server only: Use **Set Password** to change the database password. If you change a database password to match a system password, the password must satisfy the requirements of both passwords. For example, if you use the Oracle database platform, the password must meet the special character requirements that Oracle supports.

Procedure

1. In the Users application, access the user whose password you want to change.
2. Click **Set Password**.
3. In the New Password section, generate a random password or specify a password.
4. To send an Email message with the password to the user, select the **Email Password to User** check box. This check box is read-only and selected when the **Always Email Generated Passwords to Users (Never Display on Screen)** check box in the Security Controls window is selected. This check box is editable when the **Allow Generated Passwords to Be Displayed on Screen** option in the Security Controls window is selected.
5. To have the user be able to change their password upon initial login, select the **Password Should Expire After First Login** check box.
6. To synchronize this password with a database password, select the **Also Change Database Password to This Password** check box. The check box is read-only when the user ID does not have a database user ID. The check box is editable when the user ID has a database user ID that was created in the system.
7. Click **OK**.

What to do next

If you change a password for a default user (such as a system user for self-registration), then you also must change the associated property. These properties include `mxe.adminPasswd` and `mxe.system.regpassword`. You use the System Properties application to change these properties.

Related concepts:

“Passwords for users” on page 189

For security purposes, users must define a password before logging in.

“Password hints for users” on page 190

Password hints help users remember the passwords that they created and they require users to confirm their identity before an administrator resets passwords. Password hints also help users when they reset passwords.

Specifying password hints for users

You can specify password hints to check if the person requesting a new password is the user.

About this task

When you use an application server for authentication, the directory manages user passwords.

Procedure

1. In the Users application, open record for the user whose password hint you want to define.
2. From the **Select Action** menu, select **Set Password Hint**.
3. In the Set Password hint window, define the password hint question.
4. Type the answer to the question.
5. Click **OK**.

What to do next

Users also can set their own password hint by selecting **Password Information**.

Related concepts:

“Passwords for users” on page 189

For security purposes, users must define a password before logging in.

“Password hints for users” on page 190

Password hints help users remember the passwords that they created and they require users to confirm their identity before an administrator resets passwords. Password hints also help users when they reset passwords.

Specifying security groups for users

To grant users the privileges that are associated with a security group, you can specify a security group for that user.

About this task

By default, when you use an application server for authentication, the directory manages user creation. You can set properties to let user creation be performed directly in the system. The settings of these properties result in certain features being enabled or disabled in the system.

Procedure

1. In the Users application, select the user to whom you would like to assign security groups and select the **Groups** tab.
2. In the Groups section, click **Select Groups**.
3. Complete one of the following steps:
 - Select the check box next to the security group that you want to select.
 - To select all of the groups, select the **Group** check box in the table heading row.
4. Click **OK**.

Specifying security profiles for users

When you specify a security profile for multiple users, you assign the users to groups and set other security attributes. The groups and security attributes include the default insert site, the use default insert site as a display filter, the storeroom site, and the default storeroom.

Procedure

1. In the Users application, select one or more users.
2. From the **Select Action** menu, select **Set Security Profile**. If you selected multiple users, the **User Count** field displays the number of users.
3. In the Set Security Profile window, select whether you want to add, remove, or replace groups.
4. Click **New Row** and specify a group.
5. In the Update User Defaults section, select the **Edit** check box for each default that you want to specify, and specify a value for each default.
6. Click **OK**.

Results

The number of records that you updated is not always the same as the number in the **User Count** field. A record is not updated unless the change affects the current security profile. For example, if the security profile for a user contains groups A and B and you add group A, then the record for the user is not updated.

Related concepts:

“Security profiles for users” on page 191

You can specify the security profile for a user. You can assign the user to groups and set other security attributes. These attributes include default insert sites, storeroom sites, and default storerooms.

Specifying security profiles for multiple users

When you specify a security profile for multiple users, you assign the users to groups and set other security attributes. These groups and security attributes include the default insert site, the use default insert site as a display filter, the storeroom site, and the default storeroom.

Procedure

1. In the Users application on the **List** tab, select the user records for which you want to specify security profiles.
2. From the **Select Action** menu, select **Set Security Profile**. In the Set Security Profile window, the **User Count** field displays the number of users that you selected.
3. In the Set Security Profile window, select whether you want to add, remove, or replace groups.
4. Click **New Row** and specify a group.
5. In the Update User Defaults section, select the **Edit** check box above the default settings that you want to specify.
6. Specify settings for default insert site, use default insert site as a display filter, storeroom site for self-service requisitions, and default storeroom for self-service requisitions.
7. Click **OK**.

Results

The number of records that you updated is not always the same as the number in the **User Count** field. A record is not updated unless the change affects the current security profile. For example, if the security profile for a user contains groups A and B and you add group A, then the record for the user is not updated.

Related concepts:

“Security profiles for users” on page 191

You can specify the security profile for a user. You can assign the user to groups and set other security attributes. These attributes include default insert sites, storeroom sites, and default storerooms.

Granting user access to Oracle and Structured Query Language server databases

You can grant a user authorization to read, insert, update, and delete specific objects that define a set of fields and business rules, and that update one or more database tables.

Before you begin

Before you can grant database access, you must perform additional steps to grant authorization to database users. To provide a user complete database authorization, use the tools and procedures of your database platform. If you use the IBM DB2 database platform, you must use DB2 tools to provide database access rights. If you use an application server for authentication and user management, you cannot use change database access.

Procedure

1. In the Users application, access the user who needs database access.
2. From the **Select Action** menu, select **Database Access**.
3. In the Database Access window, type an ID for the user. The ID must meet the requirements of the database platform and can be different from the user ID.
4. In the **Database Password** field and **Confirm Password** fields, type a database password. This password can be different than the user password.
5. To grant a user access to specific tables or objects in the database, complete the following steps:
 - a. Click **New Row**.
 - b. Specify the name of the object. The name of the entity populates the **Entity Name** field.
 - c. To define user rights to the tables associated with the object, select the **Read**, **Insert**, **Update** and **Delete** check boxes.
6. Click **OK**.

Changing user access to Oracle and Structured Query Language server databases

To manage database access for users, you can change their existing access. You can remove rights to objects, and add, delete, or change existing rights.

Before you begin

If you use the IBM DB2 database platform, you must use IBM DB2 tools to provide database access rights. If you use an application server for authentication and user management, you cannot change database access.

Procedure

1. In the Users application, open the user for whom you want to change database access.
2. From the **Select Action** menu, select **Database Access**.

3. In the Database Access window, remove specific object rights; and add, delete, or change the rows.
4. Click **OK**.

Related concepts:

“Database users” on page 186

Database users are granted access to read, insert, update, and delete specific objects that define a set of fields and business rules. These objects can also update one or more database tables.

“Database access for users” on page 191

By default, users do not have authorization to access the database. However, in some instances, users need access to the database. For example, a user might need database authorization to view tables and columns to create reports.

Removing user access to Oracle and Structured Query Language server databases

You can remove access for a user to read, insert, update, and delete specific objects. These objects define a set of fields and business rules, and update one or more database tables.

Before you begin

If you use the IBM DB2 database platform, you must use DB2 tools to provide database access rights. If you use an application server for authentication and user management, you cannot use change database access.

Procedure

1. In the Users application, access the user for whom you want to delete database access.
2. From the **Select Action** menu, select **Database Access**.
3. In the Database Access window, click **Drop Database User**.
4. Click **OK**.

Logging out and blocking users

You can manage the ability of a user to log in. For users who are currently logged in, you can either log them out or block them.

Procedure

1. From the **Select Action** menu in the Users applications, select **Manage Sessions**.
2. Click the **Current Sessions** tab. The information in the Current Sessions table windows is read-only. You can sort and filter the information in each column, and also download it.
3. In the Current Sessions table window, click **View Details** for the user who you want to manage.
4. Log out the user or log out and block the user.

Enabling login tracking

You use the **Security Controls** action in the Security Groups and Users applications to enable login tracking. Login tracking enhances security by limiting the number of incorrect passwords a user can enter when attempting to sign in.

Before you begin

You must enable login tracking to control the number of login attempts allowed for a user. If you set the number of login attempts without enabling login tracking, the asset management system does block the user when the number of maximum number of login attempts is exceeded.

About this task

When you enable login tracking, the asset management system logs all sign in attempts, successful and unsuccessful. You can specify the maximum number of unsuccessful sign in attempts. If a user exceeds the maximum number of unsuccessful attempts, the status of the user record is changed to **BLOCKED**. The user is prevented from logging in until an administrator uses the **Change Status** action to set the status back to **ACTIVE**.

Procedure

1. From the Select Action menu in the Security Groups or Users application, select **Security Controls**.
2. In the Security Controls window, specify whether you want to enable the login tracking:

Option	Enable Login Tracking
Enabled login tracking	Selected
Disabled login tracking	Cleared

3. Enter values in the following fields:

Option	Description
Login Attempts Allowed	The number of times a user can incorrectly enter their user name or password before being blocked from the asset management system.
Password lasts this Number of Days	The number of days a user password is valid. If you do not want the password to expire, leave this field empty.
Days Before Password Expires to Warn User	The number of days before the user receives an expiration warning message.
Days Before Previously Used Password can be Reused	The number of days before a user can use a previously used password. If this value is 0, the asset management system does not check for password reuse.

4. Click **OK**.

Setting user defaults

You use the Security Controls action to specify the defaults for user records. You can access the Security Controls action from either the Security Groups application or the Users application.

About this task

You can specify the following defaults for user records:

- Default security group for new users - New users are automatically assigned to a security group. The default group defines their security permissions until they are assigned to additional groups. The default group is MAXDEFLTREG. The permissions for this group are limited to access to the Start Center. Users can change their own passwords.
- Default status for new user records - The default status is NEWREG. The NEWREG status allows you to search for new user records. You can also route records into a workflow process.
- Group for all users - The default security group for global permissions is MAXEVERYONE. When you select **Group for All Users** in Security Controls, the user is added to MAXEVERYONE.
- Electronic signature dialog - When you select **Display User ID in the Electronic Signature Dialog** in Security Controls, the system displays the user ID in the window, and prompts the user to enter a password.

Your implementation might use an application server to authenticate with an external directory by means of the Lightweight Directory Access Protocol (LDAP). In this case, you do not use the system to perform some functions. These functions include:

- Self registration - This function is not supported in conjunction with an external directory.
- Setting or changing passwords and password hints - All password-related functions are managed by the directory.

By default, when you use an application server for authentication, the directory manages user and group creation. You can set properties to allow user and group creation to be performed directly in the system. The settings of these properties result in certain features being enabled or disabled in the system.

Procedure

1. From the **Select Action** menu, select **Security Controls**.
2. In the User Defaults section, specify the following defaults:
 - a. In the **Default Group for New Users** field, type the name of the group, or click **Detail Menu**.
 - b. In the **Initial Self-Registered User Status** field, type a user status.
 - c. In the **Group for All Users** field, specify the group for global permissions. The default is MAXEVERYONE.
 - d. Select the **Display User ID in the Electronic Signature Dialog** check box to display the user ID in the window when the system prompts users to enter their passwords. If you implement electronic signatures, you must enable login tracking.
3. Click **OK**.

Copying users

If you are creating a user record and want to use settings that are similar to those in an existing record, you can copy the existing one.

Before you begin

When you copy a user record, the user settings, the purchasing general ledger accounts, and the group membership information is copied. You can change the values on the new user record.

About this task

By default, when you use an application server for authentication, the directory manages user creation. You can set properties to let user creation be performed directly in the system. The settings of these properties result in certain features being enabled or disabled in the system.

Procedure

1. In the Users application, open the user record that you want to copy.
2. From the **Select Action** menu, choose **Duplicate User**.
3. Optional: In the **User** field, type a unique user ID. If you use autonumbering, the **User** field is already populated.
4. Move the cursor to the **Type** field.
 - a. Optional: If no matching person ID is found, select one of the following options to create a person ID:
 - To create a person record based on the information in the Personal section of the user record that you are copying, click **Yes** in the System Message window.
 - To select an existing person record, click **No** and select the person record in the **Person** field.
 - To create a person record, click **No**, and in the **Person** field, select **Detail Menu** and then select **Go to People**.
5. Specify a type and person.
6. Define a password.
7. Optional: Enter additional information about the user in the Personal section, User Setting section, and Purchasing section.
8. Save your changes.

Deleting users

To manage employee information, you can delete user records that you no longer need.

Before you begin

User records cannot be deleted if the user ID is a system account or is specified as the **Run As User** for an active cron task instance. If you delete users in the directory server, the VMMSYNC cron task does not delete the users from the system tables.

About this task

When you delete a user and login tracking is not enabled, all evidence of the user is removed from the database. The user name can then be reused. If login tracking is enabled, the user name is retained in the MAXUSER table with a status of Deleted and the user name cannot be reused. You can configure login tracking in the Security Controls window

Procedure

1. From the **Select Action** menu, select **Delete User**.
2. In the Confirmation window, click **Yes**.

Results

The user record is deleted and that person can no longer login.

Deleting security groups from user profiles

To manage the security privileges associated with user records, you can delete a security group from a user profile.

Before you begin

You can delete a security group from a user profile once the user is assigned to another security group.

About this task

When you delete a security group from a user profile, the association between the security group and the user is removed. The group then remains in the database.

Procedure

1. In the Users application, open the user record whose groups you want to change.
2. Click the **Groups** tab.
3. To delete a group, click **Mark Row for Delete**.
4. Save your changes.

Related concepts:

“Security profiles for users” on page 191

You can specify the security profile for a user. You can assign the user to groups and set other security attributes. These attributes include default insert sites, storeroom sites, and default storerooms.

Chapter 8. Managing communication templates

You use communication templates to standardize frequently used email communications, which are also known as notifications. There are several ways you can use communication templates, such as with the workflow or escalation process. You can also create and send email communications from the ticket applications by using standardized information from communication templates.

Communications template overview

You can create and manage communication templates that can be used to standardize frequently used e-mail communications (which are also known as notifications).

Communication templates and escalations

You can use escalations to monitor time-sensitive records and key performance indicators (KPIs). When you create escalation records, you can specify that email notifications are generated when a record reaches the defined escalation point.

You can create each notification in the Escalations application, or you can create communication templates for frequently generated notifications. Notifications are sent when records are found that meet the conditions that an escalation point defines.

In the Escalations application, you can create a template-based communication. A template-based communication uses all the features that are available in a communication template, including attachments. The default values for role or recipient, subject, and message are derived from the template. You cannot change these values in the Escalations application.

For example, a service desk agent does not complete assignments within the specified time span of six hours. The assignment is escalated to the supervisor and the supervisor receives an email communication.

Communication templates and the service desk

You can use communication templates to create and send email messages from the service desk applications. These applications include the Service Requests application, the Incidents application, and the Problems application.

When you create communication templates for the service desk applications, you ensure that communications with service desk customers contain standardized information.

Communication templates and workflow

You can design a workflow process to generate notifications about the progress of a specific record. Notifications can be made through email or through pager, providing that your paging system supports email.

When you design a workflow process that includes email notifications, you can create the notification, and apply a communication template and change or complete the notification.

For workflow processes, you create templates that use roles as a recipient. You can add one or more roles. You can add more than one category of recipient, such as persons, person groups, or email addresses. In the Workflow Designer application, you can create a template-based communication. A template-based communication uses all the features available in a communication template, including attachments. The default values for role or recipient, subject, and message are derived from the template. You cannot change these values in the Workflow Designer application.

You can create a template for purchase requisition approvals or rejections, which can be sent as the request flows through the workflow process. You create a workflow process for purchase requests. A user submits a request for a notebook. The request enters the workflow process and is approved by the immediate supervisor. After the purchase requisition is approved by finance, the status is set to approved and the user is notified of the approval.

Substitution variables for communication templates

In the Communication Templates application, you use substitution variables in the **Subject** field and in the **Message** field. When you use substitution variables in a communication template that is used to create a notification, the substitution variables are replaced with the corresponding values in the record that generates the notification.

When you use substitution variables in the **Message** field, you type the text in the **Subject** field or in the **Message** field. You then add a space and a colon before the substitution variable to format the output correctly. If more text or other variables follow, you must insert a space after any variables in the **Subject** field or the **Message** field.

You can also use dot notation with relationships in substitution variables. For example, rel1. rel2. fieldname.

Your Incident #:TICKETID was opened on :REPORTDATE . The person assigned to work on your issue is :OWNER . You will be contacted on or before :TARGETSTART .

Predefined communication templates

In the Communication Templates application, there are predefined communication templates that you can use to create notifications. You can use the predefined templates with the Workflow application, the Service Requests application, and the Incidents application. You can also use the predefined templates with the system database or with the demo database (MAXDEMO).

Templates for the system database

You can customize templates for the system database for your business needs. Because there is an associated escalation or workflow process that refers to these templates, you cannot delete them.

Templates for the demo database

You can change or delete templates for the demo database (MAXDEMO). You can use the demo database communication templates in your test environment to gain experience adding and managing templates.

The email Listeners application uses default templates for error notifications. If an error is encountered while staging inbound records, an error notification is sent to the email administrator. The type of error determines the error notification that is sent to the email administrator.

Recipients of communication templates

In the Communication Templates application, you can add recipients to communication templates. You can choose from four types of recipients: roles, persons, person groups, and email addresses. You can add one or more recipients from each category, and you can add more than one type of recipient.

If you are creating a communication template for use with a workflow process or escalation process, you must add at least one recipient. If you are creating a template for use with the Ticket application or the Work Order application, you do not have to add a recipient. If you do not add a recipient, users who apply the template to a record are notified that an email is not sent, because there are no recipients.

Attachments for communication templates

In the Communication Templates application, you can attach many types of document files to communication templates. These document files include text files, images, spreadsheets, videos, web pages, and document folders. When you create a communication based on a template, the attachments are always sent with the communication.

Communication logs

In the Communication Templates application, a communication log lists the inbound and outbound communications for a record, such as a ticket or work order. For outbound communications that are generated with the **Create Communication** action, the communication log entry contains the details from the email message and any attachments.

Working with communication templates

You can create and manage communication templates that can be used to standardize frequently used e-mail communications also known as notifications.

Creating communication templates

You can create communication templates to standardize frequently used email communications. You can also use communication templates to create email notifications to use with the automated workflow and escalation processes.

Procedure

1. In the Communication Templates application, click **New Communication Template**. A new communication template opens with an inactive status.
2. Optional: If the **Template** field is blank, provide a name or identifier.
3. In the **Applies To** field, specify a value.
4. In the **Accessible From** field, specify from where users can access the template:

Option	Description
ALL	For the template to be available: • From the Create Communication action in other applications • For use with workflow and escalation processes
APPS	For the template to be available from the Create Communication action in other applications, except for the Escalations application and the Workflow application
ESCALATION	For the template to be available only with the escalation function
WORKFLOW	For the template to be available only with the workflow function

5. Optional: Create an entry in the communication log and attach files to the communication template.
6. Complete the following details for the communication template:
 - a. Specify the email address from which the communication template is sent.
 - b. Optional: If the recipient must reply to an email address other than the address of the sender, provide an email address in the **Reply to** field.
 - c. In the **Subject** field and **Message** field, specify substitution variables.
 - d. Optional: Add a recipient on the **Recipients** tab. If you are creating a communication template for a workflow or escalation process, you must add at least one recipient.
 - e. Optional: If you are creating a communication template for a workflow process and you want the last memo to display first in the list of memos, type `:wfassignment.lastmemo` as part of the description for the template.
 - f. Optional: If you are creating a communication template for a workflow process and you need memos to be available immediately, type `:wfassignment.currentmemo` as part of the description for the template.
7. Save your changes.

Related concepts:

“Communication templates and escalations” on page 209

You can use escalations to monitor time-sensitive records and key performance indicators (KPIs). When you create escalation records, you can specify that email notifications are generated when a record reaches the defined escalation point.

“Communication templates and the service desk” on page 209

You can use communication templates to create and send email messages from the service desk applications. These applications include the Service Requests application, the Incidents application, and the Problems application.

“Communication templates and workflow” on page 209

You can design a workflow process to generate notifications about the progress of a specific record. Notifications can be made through email or through pager, providing that your paging system supports email.

“Predefined communication templates” on page 210

In the Communication Templates application, there are predefined communication templates that you can use to create notifications. You can use the predefined templates with the Workflow application, the Service Requests application, and the Incidents application. You can also use the predefined templates with the system database or with the demo database (MAXDEMO).

“Substitution variables for communication templates” on page 210

In the Communication Templates application, you use substitution variables in the **Subject** field and in the **Message** field. When you use substitution variables in a communication template that is used to create a notification, the substitution variables are replaced with the corresponding values in the record that generates the notification.

Adding email addresses as communication template recipients

When you create a communication template, you can use email addresses as a recipient. You can add one or more email addresses, and you can add more than one category of recipient such as persons, person groups, or roles.

Procedure

1. In the Communication Templates application, open or create a communication template.
2. Click the **Recipient** tab.
3. Click **Show Table** to expand the email table window.
4. Click **New Row** to add a recipient.
5. Select whether the recipient receives the communication directly, is copied, or is blind copied.
6. Save your changes.

Related concepts:

“Recipients of communication templates” on page 211

In the Communication Templates application, you can add recipients to communication templates. You can choose from four types of recipients: roles, persons, person groups, and email addresses. You can add one or more recipients from each category, and you can add more than one type of recipient.

Adding person groups as communication template recipients

When you create a communication template, you can use person groups as a recipient. You can add one or more person groups. You can also add more than one category of recipient such as email addresses, persons, or roles.

About this task

The email communication is sent to the first available person according to the calendar and shift. If no one is available, the email communication is sent to the default person in the person group.

Procedure

1. In the Communication Templates application, open or create a communication template.
2. On the **Recipient** tab, click **Show Table** to expand the Person table window.
3. Click **Select Groups** to add multiple recipients. In the Select Person Groups window, choose the groups that you want to add and click **OK**. For each person in the person group, you can select any of the following options:
 - **To** - the person receives the email communication directly.
 - **cc** - the person is copied on the email communication.
 - **bcc** - the person is blind copied on the email communication.
4. Optional: If you do not want the communication sent to every person in the group, clear the **Broadcast** check box.
5. Save your changes.

Related concepts:

“Recipients of communication templates” on page 211

In the Communication Templates application, you can add recipients to communication templates. You can choose from four types of recipients: roles, persons, person groups, and email addresses. You can add one or more recipients from each category, and you can add more than one type of recipient.

Adding persons as communication template recipients

When you create a communication template, you can use persons as a recipient. You can add one or more persons. You can also add more than one category of recipient such as person groups, email addresses, or roles.

Procedure

1. In the Communication Templates application, open or create a communication template.
2. Click the **Recipient** tab.
3. Click **Show Table** to expand the Person table window.
4. Select one of the following options:

Option	Description
New Row	To add a single recipient, in the Person field, specify a value.
Select People	To add multiple recipients, in the Select People window, select the people that you want to add and click OK .

5. Select whether the recipient receives the communication directly, is copied, or is blind copied.
6. Save your changes.

What to do next

On the **Communication Template** tab, you can view the recipient that you added.

Related concepts:

“Recipients of communication templates” on page 211

In the Communication Templates application, you can add recipients to communication templates. You can choose from four types of recipients: roles, persons, person groups, and email addresses. You can add one or more recipients from each category, and you can add more than one type of recipient.

Adding roles as communication template recipients

When you create a communication template, you can use roles as a recipient. You can add one or more roles. You can also add more than one category of recipient such as persons, person groups, or email addresses.

Procedure

1. In the Communication Templates application, open or create a communication template.
2. Click the **Recipient** tab.
3. Click **Show Table** to expand the Role table window.
4. Choose one of the following options:

Option	Description
New Row	To add a single recipient. In the Role field, specify a role.
Select Roles	To add multiple recipients. In the Select Roles window, select the roles that you want to add and click OK .

5. Choose whether the recipient receives the communication directly, is copied, or is blind copied.
6. Save your changes.

What to do next

On the **Communication Template** tab, you can view the roles that you added.

Related concepts:

“Recipients of communication templates” on page 211

In the Communication Templates application, you can add recipients to communication templates. You can choose from four types of recipients: roles, persons, person groups, and email addresses. You can add one or more recipients from each category, and you can add more than one type of recipient.

Attaching documents to communication templates

You can send additional information in an email communication using a communication template. You can attach document folders, files, and web pages to email communications.

Attaching document folders to communication templates

To send additional information with a communication template, you can attach document folders. When you create a communication based on a template, the document folders are always sent with the communication.

About this task

The business object to which the template applies determines which document folders you can view in the Folders table window. These folders are defined in the originating application. For example, if you create a communication template for incidents, the system lists any document folders that have been defined in the Incidents application. The system also links any folders linked in the database to the Incidents application.

Procedure

1. In the Communication Templates application, open or create a communication template.
2. Click the **Attachment Folders** tab.
3. Select the **Send with Communication** check box for the document folders that you want to attach to the template.
4. Save your changes.

Related concepts:

“Attachments for communication templates” on page 211

In the Communication Templates application, you can attach many types of document files to communication templates. These document files include text files, images, spreadsheets, videos, web pages, and document folders. When you create a communication based on a template, the attachments are always sent with the communication.

Attaching files to communication templates

To send additional information with a communication template, you can attach many types of files, including text files, images, spreadsheets, and videos. When you create a communication based on a template that has attachments, the attachments are always sent with the communication.

Procedure

1. In the Communication Templates application, open or create a communication template.
2. On the **Communication Template** tab, click **Attachments** and select whether to attach files or attach files from the library.
3. Save your changes.

Related concepts:

“Attachments for communication templates” on page 211

In the Communication Templates application, you can attach many types of document files to communication templates. These document files include text files, images, spreadsheets, videos, web pages, and document folders. When you create a communication based on a template, the attachments are always sent with the communication.

Attaching web pages to communication templates

To send additional information with a communication template, you can attach web pages. When you create a communication based on this template, the attachments are always sent with the communication.

Procedure

1. In the Communication Templates application, open or create a communication template.
2. On the **Communication Template** tab, click **Attachments**.

3. To attach a Web page, select **Add New Attachments > Add New Web Page**.
4. Save your changes.

Related concepts:

“Attachments for communication templates” on page 211

In the Communication Templates application, you can attach many types of document files to communication templates. These document files include text files, images, spreadsheets, videos, web pages, and document folders. When you create a communication based on a template, the attachments are always sent with the communication.

Linking records to communication templates

You can insert a hyperlink to a record within the body of your communication template message that includes the application name and record ID.

Procedure

Insert the following link into the body of your message:

```
http://:HOSTNAME/maximo/ui/maximo.jsp?event=loadapp&value=:APP&uniqueid=:OWNERID
```

where *HOSTNAME* is the name or IP address of the MXServer

where *OWNERID* is the name of the owner

Results

When a notification is generated, the application name and record ID appear in the message as a hyperlink that leads directly to the record.

To access the record, the email recipient must be a registered user with a security permission in the specified application.

Copying communication templates

You can create a template that is based on an existing template.

About this task

When you copy a communication template, the copy has the same information as the existing template, except for a new template ID and an inactive status.

Procedure

1. In the Communication Templates application, display the template that you want to copy.
2. From the **Select Action** menu, select **Duplicate Template**.
3. Optional: If the **Template** field is blank, specify a value.
4. Save your changes.

Related concepts:

“Communication templates and escalations” on page 209

You can use escalations to monitor time-sensitive records and key performance indicators (KPIs). When you create escalation records, you can specify that email notifications are generated when a record reaches the defined escalation point.

“Communication templates and the service desk” on page 209

You can use communication templates to create and send email messages from the service desk applications. These applications include the Service Requests application, the Incidents application, and the Problems application.

“Communication templates and workflow” on page 209

You can design a workflow process to generate notifications about the progress of a specific record. Notifications can be made through email or through pager, providing that your paging system supports email.

“Predefined communication templates” on page 210

In the Communication Templates application, there are predefined communication templates that you can use to create notifications. You can use the predefined templates with the Workflow application, the Service Requests application, and the Incidents application. You can also use the predefined templates with the system database or with the demo database (MAXDEMO).

“Substitution variables for communication templates” on page 210

In the Communication Templates application, you use substitution variables in the **Subject** field and in the **Message** field. When you use substitution variables in a communication template that is used to create a notification, the substitution variables are replaced with the corresponding values in the record that generates the notification.

Changing communication templates

Once you create a communication template, you can change it to reflect any changes in your business needs.

Before you begin

You do not need to deactivate the template before you change it.

Procedure

1. In the Communication Template application, open the template that you want to change.
2. Change the information in the appropriate fields.
3. Save your changes.

Related concepts:

“Communication templates and escalations” on page 209

You can use escalations to monitor time-sensitive records and key performance indicators (KPIs). When you create escalation records, you can specify that email notifications are generated when a record reaches the defined escalation point.

“Communication templates and the service desk” on page 209

You can use communication templates to create and send email messages from the service desk applications. These applications include the Service Requests application, the Incidents application, and the Problems application.

“Communication templates and workflow” on page 209

You can design a workflow process to generate notifications about the progress of a specific record. Notifications can be made through email or through pager, providing that your paging system supports email.

“Predefined communication templates” on page 210

In the Communication Templates application, there are predefined communication templates that you can use to create notifications. You can use the predefined templates with the Workflow application, the Service Requests application, and the Incidents application. You can also use the predefined templates with the system database or with the demo database (MAXDEMO).

“Substitution variables for communication templates” on page 210

In the Communication Templates application, you use substitution variables in the **Subject** field and in the **Message** field. When you use substitution variables in a communication template that is used to create a notification, the substitution variables are replaced with the corresponding values in the record that generates the notification.

Deleting communication templates

You can delete a communication template that is no longer being used.

Before you begin

You can delete a communication template only if it is not referenced by a workflow process or escalation process.

Procedure

1. In the Communication Templates application, open the template that you want to delete.
2. From the **Select Action** menu, select **Delete Template**.
3. Click **Yes**.

Related concepts:

“Communication templates and escalations” on page 209

You can use escalations to monitor time-sensitive records and key performance indicators (KPIs). When you create escalation records, you can specify that email notifications are generated when a record reaches the defined escalation point.

“Communication templates and the service desk” on page 209

You can use communication templates to create and send email messages from the service desk applications. These applications include the Service Requests application, the Incidents application, and the Problems application.

“Communication templates and workflow” on page 209

You can design a workflow process to generate notifications about the progress of a specific record. Notifications can be made through email or through pager, providing that your paging system supports email.

“Predefined communication templates” on page 210

In the Communication Templates application, there are predefined communication templates that you can use to create notifications. You can use the predefined templates with the Workflow application, the Service Requests application, and the Incidents application. You can also use the predefined templates with the system database or with the demo database (MAXDEMO).

“Substitution variables for communication templates” on page 210

In the Communication Templates application, you use substitution variables in the **Subject** field and in the **Message** field. When you use substitution variables in a communication template that is used to create a notification, the substitution variables are replaced with the corresponding values in the record that generates the notification.

Changing the status of communication templates

You can manage the use of communication templates by changing the status to either active or inactive. An active status indicates that the template is ready for approval or use. An inactive status indicates that you no longer want to use the template.

Before you begin

If the template is used for an active escalation or workflow process, you cannot change the status of a template to inactive. Only templates that you set to active can be applied by other users to a ticket record.

About this task

The default status of a template is inactive.

Procedure

1. In the Communication Templates application, open a communication template.
2. Click **Change Status**.
3. In the **Status** field, select a status. You can view the date and time of the status change in the **Status Date** field.
4. Click **OK**.

Related concepts:

“Communication templates and escalations” on page 209

You can use escalations to monitor time-sensitive records and key performance indicators (KPIs). When you create escalation records, you can specify that email notifications are generated when a record reaches the defined escalation point.

“Communication templates and the service desk” on page 209

You can use communication templates to create and send email messages from the service desk applications. These applications include the Service Requests application, the Incidents application, and the Problems application.

“Communication templates and workflow” on page 209

You can design a workflow process to generate notifications about the progress of a specific record. Notifications can be made through email or through pager, providing that your paging system supports email.

“Predefined communication templates” on page 210

In the Communication Templates application, there are predefined communication templates that you can use to create notifications. You can use the predefined templates with the Workflow application, the Service Requests application, and the Incidents application. You can also use the predefined templates with the system database or with the demo database (MAXDEMO).

“Substitution variables for communication templates” on page 210

In the Communication Templates application, you use substitution variables in the **Subject** field and in the **Message** field. When you use substitution variables in a communication template that is used to create a notification, the substitution variables are replaced with the corresponding values in the record that generates the notification.

Chapter 9. Managing escalations

You use escalations to ensure that critical tasks, such as those defined in service level agreements, are completed on time. You can also use escalations for events that notify you before contracts expire, or to change the status or owner of a record. You can use escalations with any application.

Escalations overview

An escalation is a mechanism for monitoring time-sensitive records. Escalations can take actions or send notifications when a record reaches a defined escalation point. Your workflow administrator can specify that a task assignment has a time limit. If the assignment is not completed within the specified time limit, an escalation can be triggered for the record.

Escalation engine

The product contains an escalation engine that runs the escalations. Escalations help you ensure that critical tasks, such as those defined in service level agreements, are completed on time.

The escalation engine performs the following tasks:

- Drives the escalation process
- Uses the cron task function
- Tests all active escalation definitions at a set schedule
- Retrieves escalation definitions from the database and constructs appropriate SQL statements
- Runs SQL statements against target objects for the escalation
- Retrieves records, and performs actions and notifications associated with the escalation definitions

Escalation logs

You can monitor the execution of escalations by the escalation engine. To do so, you can configure logging and examine the log files for log statements related to the escalation engine.

To enable escalation engine logging, you use the Logging application. For each of the loggers set the wanted log level. If the logger does not exist, create your own logger. Verify that the loggers are all active. Associate an appender with a logger to ensure that log statements send to a product log file. When you save your changes, run the **Apply Settings** action from the select action menu of the Logging application for the settings to take effect immediately.

The following table lists the escalation loggers, a description of the logger, and the log level to set each logger to.

Logger	Description	Log level
crontask	Root logger that generates log statements for all cron tasks	DEBUG

Logger	Description	Log level
ESCALATION	Child logger that inherits from cron task root logger and generates log statements specifically for the escalation engine	DEBUG
sql	Root logger that generates log statements of SQL statements run in the application server	INFO
COMMTEMPLATE	Child logger that inherits from service root logger and generates log statements specifically for communication templates and notifications events	INFO
mail	Root logger that generates log statements for communication with the mail server when sending notifications.	DEBUG

Structured Query Language Expression Builder

Use the Structured Query Language (SQL) Expression Builder to easily specify conditions to which an escalation applies. You also can build SQL expressions manually.

The SQL Expression Builder is a tool for building SQL expressions. This tool requires a basic understanding of SQL structure and syntax. The SQL statement evaluates to a true or to a false result at run time.

The SQL Expression Builder includes the following functions:

- Common SQL conditions
- Operators
- Mathematical functions
- Calendar lookup
- Classification lookup
- Keywords

The SQL Expression Builder also contains a relationship tree. This relationship tree allows you to drill down through the fields and related tables for the specified object or application and select a value.

Escalation points

In an escalation, you define the conditions to be met for an associated record, such as a work order, an asset, or a purchase order. When the conditions are met, the escalation is triggered. An escalation point represents the condition that must be met. You can use one or more of escalation points to define an escalation.

Activating an escalation does not trigger an escalation process. An escalation is triggered only when the escalation engine finds records that meet the criteria defined by the escalation points.

You can create the following categories of escalation points:

Elapsed time since a past event

Compares the current date and time to the specified field that represents an event in the past. You can select from a list of **DATETYPE** fields on the record. For example, a Start Date on a workflow assignment, an Actual Start date on a work order, or a Status Date on a record that includes status.

Time until a future event

Compares the current date and time to the specified field that represents an event in the future. For example, a Renewal Date on a contract, a Due Date on an invoice, or a Target Finish date on a work order.

Condition

Condition without a time measurement. If you want to trigger the actions and notifications of an escalation based on a condition that does not have a time measurement, you can specify the condition in the **Escalation Point Condition** field. You also can use the **Condition** field to specify that the escalation point is applied only to the subset of records specified by the condition.

Actions associated with escalation points:

When records meet the conditions in an escalation point, an action can be triggered. An action is an event, such as a changing status. Escalation points are the components of an escalation that represent a monitored condition or threshold, such as measuring elapsed time. To activate an escalation, you must associate at least one action with an escalation point.

Action types define categories of actions. You can use a predefined action type, such as set owner, status change, and create ticket, or you can create an action type.

Action groups are predefined sets of actions that are grouped in a specific sequence. Escalation points are associated with actions through the action group. If you associate a predefined action, an action group is created for that action, and the predefined action becomes a member of that action group. Only the action group with the escalation point is associated, not the action itself.

When you add multiple actions, you can assign a preexisting action a sequence number. The values for the description and the action type are defaulted based on the object and action that you chose.

Action groups can be created automatically when you add actions in the Escalations application, or you can add action groups in the Actions application. In the Actions application, you create actions and specify that an action is used with an escalation.

Predefined escalations

To help simplify escalation management, the Escalation application has two categories of predefined escalations: those escalations for the Maximo database and those escalations for the demonstration (Maxdemo) database.

To use the predefined escalations, you must activate them.

Predefined escalations for the database

You can change predefined escalations for the database to suit your business needs. You must not delete these predefined escalations; they are required for the escalations functionality to work.

The following table describes the predefined escalations for the database.

Table 37. Predefined escalations for the database

Escalation name	Application	Description
INVDUE	Invoice application	Changes the invoice status to the target invoice status when the due date is reached on invoices that were generated from a payment schedule
MSTRCTREFF	Master Contracts application	Changes the status of a contract to approved when the contract start date is reached
LEACTREFF	Lease Contracts application	
WARCTREFF	Warranty Contracts application	
PURCTREFF	Purchasing Contracts application	
LABCTREFF	Labor Contracts application	

Predefined escalations for Maxdemo database

The following table describes examples of the predefined escalations for the Maxdemo database.

Table 38. Examples of the predefined escalations for the Maxdemo database

Escalation name	Application	Description
Warranty expiring	• Master Contracts application • Lease Contracts application • Warranty Contracts application • Purchasing Contracts application • Labor Contracts application	Notifies the current owner of the warranty 90 days, 60 days, and 30 days before the expiration of the warranty of the asset

Table 38. Examples of the predefined escalations for the Maxdemo database (continued)

Escalation name	Application	Description
Contract renewal	• Master Contracts application • Lease Contracts application • Warranty Contracts application • Purchasing Contracts application • Labor Contracts application	Notifies the buyer of the contract 90 days, 60 days, and 30 days before the contract renewal date
SLA review date	• Service Level Agreements application	Notifies the service level agreement (SLA) administrator 30 days, 60 days, and 90 days before the SLA review date for all active SLA
Autoclose resolved tickets	• Service Requests application • Incidents application • Problems application	Automatically closes all non-historical tickets (all classes) that have been in a resolved status for more than 10 days

Escalations and service level agreements

Escalations help businesses comply with service level agreement commitments by proactively avoiding service level agreement violations.

- Each service level agreement has a one-to-one relationship with an escalation.
- Each commitment in a service level agreement maps to an escalation point in the corresponding escalation.
- After defining a service level agreement, you can define the corresponding escalation, and the service level agreement application populates escalation points (you can modify them).
- The service level agreement application contains an **Escalation** tab, providing a view into the corresponding escalation.

Communication templates and notifications

An escalation can initiate notifications when records are not acted upon in a timely manner. You can ensure that notifications are uniform in structure by basing them on communication templates. Notifications are sent out in the form of emails through your email service.

A notification includes a template ID, the role or recipient name, the subject of the notification, and the message. If you have information that is sent out repeatedly, you can create a communication template for it and attach it as a notification on an escalation.

You can create two types of notifications in the Escalations application:

- A free-form notification that uses only a few features of a communication template

- A template-based notification that can use all the features of a communication template, including attached documents

If you create a free-form notification the system generates a template ID for it. However, the system does not save the notification for reuse in the Communication Templates application. If you select an existing communication template, the system defaults the values in the **Role/Recipient**, **Subject**, and **Message** fields from the communication template you chose. You cannot change these read-only values from within the Escalations application.

Escalation record fields

An escalation is a mechanism to monitor records that can be acted upon or to send notifications when a record reaches a defined escalation point. In the Escalations application, you can create an escalation for any business object. Because all applications are associated with business objects, you can create escalations for any application.

An escalation record consists of the following fields

Object (Applies To field)

You create escalation records for a specific business object. The escalation engine retrieves records, from the business object, that meet the escalation point criteria.

SQL statement — (Condition field)

An escalation record can apply to all application records, or to a specific set of records. You can create an SQL statement that specifies records to which the escalation is applied. The conditions can apply to one or more tables associated with the object.

Organization and Site

Escalations are stored at the system level. You can create escalations for use with a specific organization or site.

Schedule

A schedule that defines how often the system checks for records that meet the criteria for the escalation. The polling interval can be seconds, minutes, hours, days, weeks, or months. You also can specify whether the interval is calendar or date based.

Calendar Organization

The organization that is associated with the escalation calendar.

Calendar

The calendar that specifies the days and times for which the escalation is valid.

Shift The shift that specifies the days and times for which the escalation is valid.

Escalation Point

Date-and time-based, or other condition criteria for when the actions or notifications specified on the escalation record are triggered. An escalation record can have one or more escalation points.

Actions

Any actions that must be taken when a record reaches the conditions of an escalation point. You define actions separately for each escalation point. You can associate multiple actions for each escalation point. You use the Actions application to define actions.

Notifications

Any notifications that the system must generate when a record reaches the conditions of an escalation point. You define notifications separately for each escalation point.

Deletion rules for escalations

In the Escalations application, you can manage escalations by deleting them.

The following rules apply when you delete escalations:

- When you delete an escalation, the actions and notifications are not deleted. Instead, the associations between the escalation points and the actions and notifications are removed. You delete actions in the Actions application and notifications in the Communication Templates application.
- In the Escalations application, you cannot delete escalations that are associated with service level agreements (SLA).
- If a service level agreement that is associated with an escalation is deleted, then the escalation is also deleted.

Working with escalations

You use escalations to automatically monitor the critical processes in your enterprise. You can also use escalations for events, such as contract expiration, a change in the status of a records, or a change in the ownership of a record.

Creating escalations

You create an escalation to monitor time-sensitive records that could require action. You can create an escalation for any business object and for any application. You can create escalations at the organization or site level.

Before you begin

Before you can activate and use an escalation record, you must define at least one escalation point and at least one action or notification for the escalation point. An *escalation point* defines the threshold at which an escalation is triggered. An *action* is the event that is triggered when an escalation point is reached.

Procedure

1. On the toolbar, click **New Escalation**. If the **Escalation** field is empty, specify a value.
2. In the **Description** field, type a description.
3. In the **Applies To** field, specify the object to which to apply the escalation.
4. Optional: Specify values in the **Organization** and **Site** fields. If you specify a value for either organization or site, you restrict the use of the escalation to either that organization or site. If you specify values for both organization and site, the escalation can be used only at that site.
5. Optional: Type an expression in the **Condition** field to indicate to which records the escalation applies. For example, if you want to escalate only task assignments that have a value specified in the **Time Limit** field, include the following text in your SQL statement: `TIMELIMIT is not null`. You can type the SQL condition manually. You can also use the Expression Builder to create the SQL statement.
6. In the **Schedule** field, click the **Set Schedule** icon to set how frequently to poll the database for records.

7. Optional: In the **Calendar Organization**, **Calendar**, and **Shift** fields, specify values to limit when the escalation is run.
8. Click **Save Escalation**.

Related concepts:

“Predefined escalations” on page 225

To help simplify escalation management, the Escalation application has two categories of predefined escalations: those escalations for the Maximo database and those escalations for the demonstration (Maxdemo) database.

“Escalation points” on page 224

In an escalation, you define the conditions to be met for an associated record, such as a work order, an asset, or a purchase order. When the conditions are met, the escalation is triggered. An escalation point represents the condition that must be met. You can use one or more of escalation points to define an escalation.

“Actions associated with escalation points” on page 225

When records meet the conditions in an escalation point, an action can be triggered. An action is an event, such as a changing status. Escalation points are the components of an escalation that represent a monitored condition or threshold, such as measuring elapsed time. To activate an escalation, you must associate at least one action with an escalation point.

Defining escalation points

You can define one or more escalation points for an escalation. For each escalation point, you can specify one or more actions or notifications.

Procedure

1. In the Escalations application, create, or display an escalation record.
2. In the Escalation Points table window, click **New Row**.
3. In the Row Details window, specify the condition that triggers the escalation point by selecting one of the following options:
 - If you are creating a time-based escalation point, specify values in the **Elapsed Time Attribute** field, in the **Elapsed Time Interval** field, and in the **Interval Unit of Measure** field. Type a positive number in the **Elapsed Time Interval** field to indicate a time period in the past. Type a negative number to indicate a time in the future.
 - If you are creating a condition-based escalation point, type an SQL statement in the **Escalation Point Condition** field to specify the condition that triggers the escalation. You can specify a value manually or click the **SQL Expression Builder** icon.
4. If you want the actions and notifications of the escalation point triggered more than once, select the **Repeat** check box.
5. Click **Save Escalation**.

What to do next

Now you define an action to take for each escalation point that you created.

Related concepts:

“Predefined escalations” on page 225

To help simplify escalation management, the Escalation application has two categories of predefined escalations: those escalations for the Maximo database and those escalations for the demonstration (Maxdemo) database.

“Escalation points” on page 224

In an escalation, you define the conditions to be met for an associated record, such as a work order, an asset, or a purchase order. When the conditions are met, the escalation is triggered. An escalation point represents the condition that must be met. You can use one or more of escalation points to define an escalation.

“Actions associated with escalation points” on page 225

When records meet the conditions in an escalation point, an action can be triggered. An action is an event, such as a changing status. Escalation points are the components of an escalation that represent a monitored condition or threshold, such as measuring elapsed time. To activate an escalation, you must associate at least one action with an escalation point.

Defining actions for escalation points:

You define at least one action or notification for each escalation point on an escalation record. You define actions separately for each escalation point. You use the Actions application to create action records.

About this task

To define actions for an escalation point, complete the following steps:

Procedure

1. In the Escalations application, create or display an escalation record.
2. In the Escalation Points table window, select the escalation point for which you want to define actions.
3. Click the **Actions** sub tab.
4. In the Actions window, click **New Row**.
5. In the **Action** field, specify a value.
6. Optional: Modify the **Sequence** field to indicate the order in which the action is performed.
7. Click **Save Escalation**.

What to do next

You must validate the escalation before activating it.

Related concepts:

“Predefined escalations” on page 225

To help simplify escalation management, the Escalation application has two categories of predefined escalations: those escalations for the Maximo database and those escalations for the demonstration (Maxdemo) database.

“Escalation points” on page 224

In an escalation, you define the conditions to be met for an associated record, such as a work order, an asset, or a purchase order. When the conditions are met, the escalation is triggered. An escalation point represents the condition that must be met. You can use one or more of escalation points to define an escalation.

“Actions associated with escalation points” on page 225

When records meet the conditions in an escalation point, an action can be triggered. An action is an event, such as a changing status. Escalation points are the components of an escalation that represent a monitored condition or threshold, such as measuring elapsed time. To activate an escalation, you must associate at least one action with an escalation point.

Defining notifications for escalation points:

You must define at least one action or notification for each escalation point on an escalation record. You can use a communication template to create a notification. You can also type the subject, message, and recipients manually. You define notifications separately for each escalation point.

Procedure

1. In the Escalations application, create, or display an escalation record.
2. In the Escalation Points table window, select the escalation point for which you want to define notifications.
3. On the **Notifications** sub tab, click **New Row**.
4. Perform one of the following steps:
 - In the **Template** field, specify a value.
 - Type entries in the **Role/Recipient** field, in the **Subject** field, and in the **Message** field.
5. Click **Save Escalation**.

Related concepts:

“Communication templates and notifications” on page 227

An escalation can initiate notifications when records are not acted upon in a timely manner. You can ensure that notifications are uniform in structure by basing them on communication templates. Notifications are sent out in the form of emails through your email service.

Validating escalations

You must validate an escalation record before you can activate it. Validation checks the Structured Query Language (SQL) statements to ensure that the SQL is valid and that the escalation engine can run it.

About this task

The validation process checks for syntax errors and ensures that escalation conditions have been defined in the table for the specified object. The product does not validate actions or notifications.

If an error is discovered in one or more SQL statements, the errors are written to the Validation Results section of the escalation record.

Procedure

1. In the Escalations application, display the record that you want to validate.
2. From the **Select Action** menu, click **Validate**.
3. If the validation fails, click the **Maximize** button to expand the Validation Results section and to view the error log. The SQL error might be in the **Condition** field or in the **Escalation Point Condition** fields. Correct the SQL statements and validate the record again.
4. Click **Save Escalation**.

What to do next

Activate the escalation.

Activating escalations

You can activate an escalation record that has at least one escalation point and one action or notification defined for each escalation point. The product polls for records meeting the criteria set by the escalation. If a match is found, the appropriate action is triggered.

Before you begin

An escalation record must be validated before it can be activated.

About this task

The product checks for records that meet the frequency criteria defined in the **Schedule** field of the escalation record. If records are found that match, and any of those records meet the conditions defined by the escalation points, then the escalation mechanism triggers the appropriate actions, notifications, or both.

Activating an escalation does not trigger an escalation process. The escalation process triggers only when the escalation engine finds records meeting the criteria defined by the escalation points.

When you activate an escalation, all fields on the record become read-only. You cannot edit an escalation record while it is active.

Procedure

1. In the Escalations application, display the record that you want to activate.
2. From the **Select Action** menu, select **Activate/Deactivate Escalation**. The product selects the **Active** check box in the record heading and creates an instance of the ESCALATION cron task.
3. Click **Save Escalation**.

Related concepts:

“Escalation engine” on page 223

The product contains an escalation engine that runs the escalations. Escalations help you ensure that critical tasks, such as those defined in service level agreements, are completed on time.

“Escalation logs” on page 223

You can monitor the execution of escalations by the escalation engine. To do so, you can configure logging and examine the log files for log statements related to the escalation engine.

“Deletion rules for escalations” on page 229

In the Escalations application, you can manage escalations by deleting them.

Modifying escalations

You can change existing escalations.

Before you begin

All fields on an activated escalation are read-only. You must deactivate the escalation before you can edit it.

About this task

You can modify the following elements of a deactivated escalation:

- You can delete one or more escalation points.
To activate an escalation, it must have at least one escalation point. When you delete an escalation point, the links to the associated actions and notifications are deleted.
- You can delete one or more actions or notifications associated with an escalation point.
To activate an escalation, it must have at least one action or notification defined for each escalation point.

Procedure

1. From the Escalations application, display the escalation that you want to edit.
2. If you have not deactivated the escalation, from the **Select Action** menu, select **Activate/Deactivate Escalation**.
3. On the **Escalations** tab, edit the information as needed. If a field has a **Detail Menu**, click it and select an option to retrieve a different value.
4. Click **Save Escalation**.
5. If you are ready to active the escalation, choose **Activate/Deactivate Escalation** from the **Select Action** menu.

Deactivating escalations

You deactivate an escalation record to modify the record or to delete it.

Procedure

1. In the Escalations application, display the record that you want to deactivate.
2. From the **Select Action** menu, click **Activate/Deactivate Escalation**.
3. Click **Save Escalation**.

What to do next

After you deactivate an escalation, you can perform the following actions:

- You can delete one or more escalation points. When you delete an escalation point, the links to the associated actions and notifications are deleted.
- You can delete one or more actions or notifications associated with an escalation point. To activate an escalation, it must have at least one action or notification defined for each escalation point.

Related concepts:

“Deletion rules for escalations” on page 229

In the Escalations application, you can manage escalations by deleting them.

Chapter 10. Configuring e-mail listeners

In the E-mail Listeners application, you can receive and process service requests and other types of tickets using e-mail messages. You can create or update tickets. You can also indicate whether the status is changed or queried based on specified criteria.

Related information:



MustGather: Maximo E-mail Listener (Opens in a new browser window or tab)

Testing connectivity between the application server and mail server

Testing connectivity between the application server and mail server allows you to check if the information you have entered in the E-mail Listener application is correct. You use the TestEmail.bat file to test connectivity.

Before you begin

If you use WebLogic Server, download the mail.jar file and the activation.jar file from a site for downloading open source Java code.

Procedure

1. Navigate to the \IBM\SMP\maximo\tools\maximo\internal folder and download the TestEmail.bat file.
2. Prepare to run the command.

Application server	Procedure
WebSphere Application Server	1. Open the TestEmail.bat file in a text editor. 2. Add the mail-impl.jar variable into the class path of the WebSphere Application Server folder by setting CLASSPATH=C:\IBM\WebSphere\AppServer\lib\mail-impl.jar; ..\applications\maximo\lib\j2ee.jar; ..\classes.
WebLogic Server	Copy the mail.jar file and activation.jar file to the \IBM\SMP\maximo\tools\java\jre\lib\text folder.

3. Open a command prompt, navigate to the IBM\SMP\Maximo\tools\maximo\internal folder, and run the **testemail hostname port email_account password protocol debug** command.
4. Check that the output contains the phrase connected to e-mail account.
5. If the output does not contain the phrase, check the Java exception that is created. For example, the connection can fail if an invalid user name or invalid host name was entered.

Results

The TestEmail.bat file has the following output:

```
**Email text: Connected to e-mail account email_account
DEBUG: setDebug: JavaMail version 1.3.2
DEBUG: getProvider() returning
javax.mail.Provider[STORE,imap.com.sun.mail.imap.IMAPStore,Sun Microsystems, Inc]
The Microsoft Exchange IMAP4 service is ready
*CAPABILITY IMAP4 IMAP4rev1 AUTH=NTLM AUTH=GSSAPI AUTH=PLAIN STARTTLS IDLE NAME
SPACE LITERAL*
```

Syntax of the testemail command

The TestEmail.bat file is used to test the connectivity between the application server and mail server.

Parameters

To test the connectivity between the application server and mail server, you run the **testemail** command with the following parameters and values: **hostname** can be a fully qualified host name, **port**=143, **email account** is the email account specified in the E-mail Listener application (not including the @companyname.com extension), **protocol**=imap, POP3, and **debug**=true.

E-mail Listeners overview

The E-mail Listeners application can poll multiple e-mail accounts to retrieve messages. Each account is checked at periodic intervals that you establish. Based on the subject line of an e-mail message or the contents of the e-mail message body, an e-mail listener can determine if the e-mail is new or is an updated service request, incident, or problem. An e-mail listener can also determine if an e-mail is a query for information on any business object.

The application supports the following features:

- Embedded and normal message attachments
- POP3 and IMAP e-mail protocols

The E-mail Listeners application cannot process encrypted or digitally signed e-mail messages.

E-mail processing

E-mail processing uses a predefined workflow process. Various steps in the workflow process create, update, or change the status of service requests, incidents, or problems. Other steps in the workflow process execute queries and return query results to the originator of the e-mail. You can customize the workflow process or create workflow processes to suit your needs.

All communications from the originators of the e-mail messages are captured in the service request, incident, or problem communication log. Similarly, generated communications sent to the originator of an e-mail message are captured in communication logs.

You can configure the appropriate log levels to generate detailed processing and error information from the system log file regarding how the E-mail Listeners application processes e-mails.

Email listeners components

The Email Listeners application works together with a CRON task and a workflow process to provide full functionality.

The following table lists and describes the Email Listeners components.

Table 39. Email Listeners components

Component	Purpose
Email Listeners application	The application used to create, modify, and delete email listener configurations.
Email listener cron task	A component that runs continuously on the application server and uses the CRON task infrastructure. This role is performed by the LSNRCRON cron task. This component encapsulates a staging process that processes inbound email through a staging table.
Workflow process	A workflow process that parses email information from the staging table and processes it according to the subject and contents of each message.

For these components to work correctly, you must first configure your mail servers and email accounts.

E-mail listeners process

Using e-mail listener definitions, the E-mail Listeners application polls the mail server for incoming messages and sends e-mail messages in response. The response e-mail messages confirm that a desired operation was performed on behalf of the user. The e-mail messages also notifies you if there was an error performing the operation.

When the mail server receives an incoming e-mail message, the following events occur:

- The mail server polls for incoming e-mail messages from the designated e-mail account.
- The E-mail Listeners application verifies that the senders of the incoming e-mail messages have person records. If the sender has a person record, the e-mail message is processed. If the sender does not have a person record, the e-mail message is ignored.
- A preprocessor determines whether the e-mail messages are new or updates to existing communications.
- The e-mail messages are staged, which includes the following events:
 - Extracting content from the e-mail messages, including attachments.
 - Storing content in staging tables and attached document tables.
 - Launching the e-mail processing workflow process.
 - If the e-mail listener was configured to use Java Messaging Service (JMS) queues, placing e-mail messages in the JMS queue indicating that a new e-mail message needs to be processed. A messaging component picks up this e-mail message from the queue and launches the workflow.

- Using the predefined workflow process, the following events occur:
 - Flagging an e-mail message as new or as an update to an existing service request, incident, or problem
 - For new e-mail messages, creating a service request, incident, or problem, and creating an associated communication log
 - For existing e-mail messages, updating the service request, incident, or problem, and creating an entry in the communication log

Predefined workflow process for e-mail listeners

You use workflow processes to create steps to guide records for your business process. The E-mail Listener application has a predefined workflow, Listener Business Process (LSNRBP), that is used with the E-mail Listeners application.

By default, the LSNRBP process is enabled and active. To meet your business needs, you can revise the process or create a process in the Workflow Designer application.

When you associate the LSNRBP workflow process to a configuration, the e-mail listener submits each record in the staging table to the LSNRBP process. The LSNRBP process performs the following functions based on the contents of the e-mail message:

- Changes the status of an existing ticket or other business object
- Updates an existing ticket
- Creates a ticket
- Queries any business object and return results

The LSNRBP process generates response e-mail messages at various points during e-mail processing. For example, a response containing the service request number is sent to the originator of an e-mail message whenever the application creates a service request. Response e-mail messages are also generated whenever an update is made to an existing ticket or when queries are performed.

E-mail listeners definitions

As your business needs change, you can change, copy, or deactivate your e-mail listener definitions.

Deactivation of definitions

In the E-mail Listeners application, you can deactivate an e-mail listener definition to help simplify e-mail management. For example, you can deactivate an e-mail listener definition to migrate to a new mail server or to perform routine system maintenance. To deactivate an e-mail listener definition, use the **Activate/Deactivate Listener** action. When you deactivate an e-mail listener definition, messages on the mail server are not monitored and service requests are not created.

Changes to definitions

You can change the attributes of an e-mail listener definition. You can change all values, except for the e-mail address, cron task name, cron task instance, and last time run. You must deactivate an e-mail listener definition before you can change it.

If you need to change the format of the e-mail address on an account, you must create an e-mail definition. For example, if you want to change `help@company.com` to `customer_service@company.com`. To create a definition, use the **Duplicate Listener** action in the E-mail Listeners application.

Duplication of definitions

You can duplicate an e-mail listener definition to simplify the management of e-mail records. When you duplicate an e-mail listener definition, you can specify new account information.

When you duplicate an e-mail listener definition, the following points apply:

- Some mail servers are case-sensitive.
- The case of the e-mail address that you specify is preserved.
- You cannot save an e-mail listener definition if it has an e-mail address with the same name and case as an existing e-mail address on the mail server.

Security settings for e-mail listeners

As you can create, update, query, and change the status of tickets, you can configure security settings for e-mail listeners. Using these settings, you can ensure that only authorized users can execute these functions using e-mail messages.

For the sender of an e-mail message, security authorizations are checked against the security configuration for the system. This check establishes the ability of the sender to run each specific function.

The person record is a basic requirement to be able to process e-mail messages. Additional processing of e-mail messages only occurs after the person record associated with the e-mail address of the sender has been located.

The following points apply to security settings for e-mail listeners:

- If a person record is active, the corresponding user record is found.
- If a person record does not exist or is inactive, the e-mail message is not processed. An error e-mail message is sent to the sender and the administrator.
- If a user record is found, the associated authorizations are applied when the E-mail Listeners application performs security checks on incoming e-mail messages.
- If a user record is not found, the Run As user of the cron task instance for the e-mail listener is used.

To specify security settings for e-mail listeners, you can use the **Select Security Settings** action in the E-mail Listeners application. The settings identify business objects supported by each e-mail listeners workflow process. The settings also identify the corresponding applications that must be used to determine security restrictions on incoming e-mail messages. To assign the appropriate authorizations to the users who send formatted e-mail messages, configure security settings you can use the Security Groups application.

Scenarios of security authorization

The following table describes the two security authorization scenarios that are supported when a user ID exists.

Table 40. Security authorization scenarios that are supported when a user ID exists

Scenario	Support
The user exists and has authorization to perform the operation specified in the e-mail message.	• The e-mail listener performs the security check based on the sender of an incoming e-mail message. • When the user record for the sender is located, the e-mail listener builds a security profile of the user to determine authorizations. • If the user has authorization to perform add, update and change status operations, the e-mail message is processed accordingly.
The user exists but does not have authorization to perform the operation specified in the e-mail message.	• The user can update or query only records that the user created.

Communication templates for e-mail listeners

To generate standardized e-mail notifications, the E-mail Listeners application uses communication templates to generate notifications that are sent to users and administrators. The types of notifications include confirmations, validation or processing errors, and system errors.

When the E-mail Listeners application successfully processes an incoming e-mail message, a confirmation notification is generated.

The following table lists the templates that are used for confirmation notifications.

Table 41. Templates for confirmation notifications

Templates	Confirmation situation	Recipients
LSNRBPCBSR	Confirms the creation of a service request based on an incoming free-form e-mail message.	Sender of original e-mail message.
LSNRBPCHST	Confirms the status change of records based on an incoming e-mail message.	Sender of original e-mail message.
LSNRBPQRY	Response e-mail message containing details of query results from an incoming e-mail message.	Sender of original query e-mail message.
LSNRBPUBSR	Confirms the update of a service request based on an incoming free-form e-mail message.	Sender of original e-mail message.
LSNRBPUOBJ	Confirms the update of a specified object based on an incoming e-mail message.	Sender of original e-mail message.
LSNRBPCOBJ	Confirms the creation of a specified object based on an incoming e-mail message.	Sender of original e-mail message.

A validation or processing error notification is generated when the E-mail Listeners application cannot process an incoming e-mail message. This error could be because of incorrect formatting or incomplete or invalid information in the incoming e-mail message. You must review the error in the log file or in the e-mail to resolve the error.

The following table lists the templates that are used for the validation or processing of error notifications.

Table 42. Templates for validation or processing error notifications

Template used for validation or processing error notification	Error situation	Corrective action	Recipients
LSNRAUTH	Validation error notification The sender does not have authorization to perform the operation specified in the incoming e-mail message.	Authorize the user for the specified application.	Sender of original e-mail message and system administrator.
LSNRBPAUTO	Processing error notification The e-mail listener cannot create an auto key for an attribute that was declared as auto key in the incoming e-mail message.	Verify that the attribute is declared as an auto key. If the attribute is declared as an auto key, check the logs.	Sender of original e-mail message.
LSNRBPDATE	Validation error notification The sender did not specify a properly formatted date or date time value for an attribute that was specified in the incoming e-mail message.	Provide a valid formatted date.	Sender of original e-mail message and system administrator.
LSNRBPINV	Processing error notification The sender specified an update operation for an existing record in the incoming e-mail, message, but the record does not exist.	Provide a valid record.	Sender of original e-mail message.

Table 42. Templates for validation or processing error notifications (continued)

Template used for validation or processing error notification	Error situation	Corrective action	Recipients
LSNRBPUACN	Processing error notification The sender specified an invalid action in the incoming e-mail message.	Provide a valid action.	Sender of original e-mail message.
LSNRBPUNOB	Processing error notification The sender specified an object in the incoming e-mail message that the E-mail Listeners application does not support.	Provide a valid object name. There are only 3 valid objects for create and update action: service requests, problems, and incidents.	Sender of original e-mail message.
LSNRFNKEY	Processing error notification The sender did not provide all of the primary keys for a record that was specified in the incoming e-mail message.	Provide input for the mandatory key fields.	Sender of original e-mail message.
LSNRFNREQ	Processing error notification The sender did not provide all of the required attributes for a record that the sender specified in the incoming e-mail message.	Provide input for all mandatory fields.	Sender of original e-mail message.
LSNRINVM	Validation error notification The incoming e-mail message contains a blank subject line or a subject line that exceeds the allowed length (this error occurs only in free-form e-mail messages).	Enter a subject or reduce the length of the subject.	Sender of original e-mail message.

Table 42. Templates for validation or processing error notifications (continued)

Template used for validation or processing error notification	Error situation	Corrective action	Recipients
LSNRNOPER	Validation error notification The sender of the incoming e-mail message does not have a corresponding person record.	Sender of the e-mail must be defined as a person.	Sender of original e-mail message and system administrator.
LSNRSECAPP	Validation error notification The sender of the incoming e-mail message does not have the requisite authorizations to perform the operation on the object specified in the e-mail message.	Provide the authorizations needed for the object in the Security Settings window of the E-mail Listener application.	Sender of original e-mail message and system administrator.
LSNRWFMT	Processing error notification The incoming e-mail message has invalid formatted content.	Format the content of the e-mail correctly.	Sender of original e-mail message and system administrator.

If a system error occurs when the E-mail Listeners application polls or processes an incoming e-mail message, a system error notification is generated.

The following table lists the templates that are used for system error notifications.

Table 43. Templates for system error notifications

Templates used for system error notifications	Error situation	Recipients
LSNRBPEX	Processing the incoming e-mail message.	Depending on the error, sender of the original e-mail message or system administrator.
LSNRBPQERR	Placing information about an incoming e-mail message into JMS queue.	System administrator
LSNRCFGERR	Using the listener configuration to connect to an e-mail account.	System administrator
LSNRCONNF	Connecting to the mail server to access the configured e-mail account.	System administrator

Table 43. Templates for system error notifications (continued)

Templates used for system error notifications	Error situation	Recipients
LSNRERROR	Processing the incoming e-mail message.	Depending on the error, sender of the original e-mail message or system administrator.
LSNRINBF	Staging the contents of the incoming e-mail message in the internal staging table.	System administrator
LSNRJMSEF	Connecting to the configured JMS queue because the queue connection factory information for the queue is incorrect.	System administrator
LSNRJMSF	Connecting to the configured JMS queue.	System administrator
LSNRJMS	Connecting to the configured JMS queue because the queue information is incorrect.	System administrator
LSNRJMMSSYN	Connecting to the configured JMS queue, and further processing of e-mail message is performed without the JMS queue.	System administrator
LSNRMAILER	Retrieving e-mail message from the configured e-mail account.	System administrator

Preprocessors for e-mail listeners

A preprocessor determines whether incoming e-mails messages are new requests for help or updates to existing service requests.

The preprocessor is a Java component of the e-mail listener that runs on the server when the e-mail listener recognizes a new e-mail message. The preprocessor parses the subject line of any incoming e-mail message to determine whether the message is a new request for help or an update to an existing service request. By default, a built-in preprocessor parses incoming e-mail messages. The default preprocessor value is: `psdi.common.emailstner.Preprocessor`. In the E-mail Listeners application, this value displays in the **Preprocessor** field on the **Listener** tab.

The preprocessor class extracts the string that is bounded by the object key delimiter characters, which identifies an existing service request. If a string is recognized, the preprocessor stores the string in the object key column in the e-mail listener staging table. If no string is recognized, the preprocessor leaves the object key column blank.

Customized preprocessors

The base preprocessor Java class implements a standard Java interface called the `LSNRPreprocessor`. Custom preprocessor implementations must include an implementation of the `LSNRPreprocessor` interface.

The preprocessor interface provided with the system includes the following public methods:

- Boolean isNewEmail (String del, String subject)
- String getObjectKey (String del, String subject)

In the custom Java class, you can implement both methods. Each method accepts two parameters: the delimiter string and the subject line string.

Table 44. Methods

Method	Description	Base preprocessor implementation	Custom implementation
isNewEmail()	Returns a Boolean value indicating whether the e-mail is for a new or for an existing ticket	Checks whether the Object Key Delimiter string occurs exactly twice in the subject line string	Might provide different logic to determine the new or existing ticket
getObjectKey()	• Returns a string that represents the Ticket ID • Returns null, if no ID is found	Extracts the substring between the first and last occurrences of the delimiter string in the subject line	Might provide different logic to determine the Ticket ID

Object key delimiters

The E-mail Listeners application uses an object key delimiter to identify the incoming e-mail message as an existing ticket.

Object key delimiters specify characters on either side of the service request key in the subject line of incoming e-mail messages. The default value in the **Object Key Delimiter** field is **##**. For example, service request number 1009 is represented as **##1009##**, with no spaces before or after the service request key.

You can change the default value, and replace the value before and after the ticket ID with other characters. There are no restrictions regarding the characters that you can use. However, the delimiter must be unique. Choose characters or symbols that are not frequently used. For example, you use **+** as the object key delimiter, and a user sends an e-mail message with the subject line: **+1003+ Having problems with printer + network**. The preprocessor cannot identify the substring because the object key delimiter symbol occurs multiple times in the subject line.

When you send e-mail messages that reference existing service requests, place the characters that you choose as the object key delimiter before and after the service request key. The preprocessor class extracts the string that is bounded by the object key delimiter characters, which identifies an existing service request. If a string is recognized, the preprocessor stores it in the object key column in the e-mail listener staging table. If no string is recognized, the preprocessor leaves the object key column blank.

Object key identifier

An object key identifier is the ID of a record. The object key identifier can be a sequence that is generated. For example, object key identifiers can be 1001, 1002, and so on.

Logging

You can use logging to create and to manage log files that contain informational, warning, or error messages regarding the processing of e-mail messages.

You use the Logging application to create and to manage log files.

Java Message Driven Bean

To enable an e-mail listener to use JMS queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation.

E-mail messages

In the E-Mail Listeners application, you can manage your e-mail messages.

Polling of mail servers for email messages

Using email listener definitions, the email Listeners application uses a dedicated cron task to poll the mail server for incoming messages. Cron tasks are Java components that you can set to run automatically and on a fixed schedule. For each email listener definition that you create, a specific instance of the cron task is created and associated with the email listener.

The scheduling value determines the polling frequency. The polling frequency is stored as an attribute of each cron task instance. Specify different schedules for different email accounts. By default, the schedule for the polling frequency is set to every 5 minutes.

The cron task polls the mail server at the set frequency and performs the following actions:

Table 45. Cron task actions

Condition	Mail server action
For new email messages on the mail server	The mail server: • Extracts the header and message body • Extracts any attachments • Moves the email to e-mail staging • Marks the email message as read
For email messages marked as read on the mail server	The mail server determines whether read email messages are deleted: • If an email message is to be deleted, it marks the email message as deleted • If an email message is not to be deleted, it keeps the email message on the server

For email messages processed from a POP-based mail account, the Email Listener determines whether an email is already processed or is new.

Queues:

A queue is an application server component that can facilitate parallel processing. Associating a queue with an e-mail listener can speed up the sequential processing of incoming e-mails, especially when they are arriving in high volume.

The E-mail Listeners application processes incoming e-mail messages in a sequential manner. The sequence of processing includes polling the mail server, staging the mail into the database, and launching workflow processing on the staged mail record. In situations where high volumes of e-mail messages must be processed efficiently, this sequential processing can be time-consuming. This means that e-mail messages are not always processed as quickly as possible. Therefore, it can be beneficial to switch the e-mail listener to a parallel processing mode. To perform this switch, you can configure a queue and associate the queue with the listener. The system uses Java Messaging Service (JMS) queues, provided by the underlying Java application server.

Once a message is placed in a queue, the message processing component can pick the message up in an asynchronous manner. In the Java application server, the processing components are called message-driven beans (MDBs). You can configure the application server to provide multiple MDBs that process multiple messages in parallel. This increases the speed at which the E-mail Listeners application processes e-mails.

Once you set up the queues, you can modify or create an e-mail listener definition to specify queue-based processing, the queue name, and the queue connection factory name.

Staging e-mail messages:

Using an e-mail listener cron task, the E-mail Listeners application stages e-mail messages to a staging table. When staging an e-mail message, the application saves all information that is required to process the e-mail message and initiate the workflow process. Cron tasks are behind-the-scene jobs that run automatically and on a fixed schedule.

A staging table stores the attributes of an incoming e-mail message, including the recipients, the sender, the subject, and the message. The staging process creates a record, and the workflow process determines how to process the record.

If an incoming e-mail message generates a ticket, details of the e-mail message are stored as the initial entry in the communication log of the ticket. Any further correspondence is stored, based on the key of the ticket. Graphics, whether they are embedded in or attached to the e-mail message, are also visible in the communication log. You can view communication logs from the **Log** tab of the Service Request application.

Errors found during staging

If the E-mail Listeners application finds any errors while staging e-mail records, the e-mail listener can write error information in the log file on the server. You use the Logging application to enable logging for the E-mail Listeners application.

In addition, you can receive error notifications if the E-mail Listeners application encounters errors during staging. To receive the error notifications, you must provide a valid e-mail address in the **Administrator E-mail** field. The notification that is sent to the e-mail address for the administrator contains a detailed description of the errors.

Status of e-mail records

Every e-mail record in the staging table has an associated status. These statuses reflect the sequence of actions that the e-mail listener performs on an e-mail record in the staging table.

The possible statuses for e-mail records are described in the following table.

Table 46. Statuses of e-mail records

Status	Description
New	• If an e-mail message is received with a subject, an e-mail record for the e-mail message is created in the staging table. The status of the e-mail record is set to new. • The e-mail record is separated into discrete components (such as to, from, cc, subject, and message). Any attachments are extracted and stored on the file system where the system server resides. • Attachments are linked to new attached documents records. The records are linked to the e-mail record in the staging table.
In process	• During the processing of a new e-mail record, the status of the e-mail record is set to in process.
Workflow	• When the workflow process is initiated, the status of the e-mail record is set to workflow. • The workflow process either creates a new service request and communication log, or updates an existing service request.
Invalid	• If an e-mail message is received without a subject, an e-mail record for the e-mail message is created in the staging table. The status of the e-mail record is set to invalid. • The e-mail record remains in the staging table. The e-mail record is not processed further. That is, the e-mail record does not enter the workflow process. • An administrator can use the resubmit function to correct an e-mail record with an invalid status. When an e-mail is resubmitted, the status is changed to new. The e-mail record is processed as a new e-mail message.

Table 46. Statuses of e-mail records (continued)

Status	Description
Error	• If an error is encountered during the staging of e-mail records, the status of the e-mail record is set to error. An error notification is sent to the administrator specified during the configuration of the e-mail listener. • If an e-mail address for the administrator was not specified and if logging is enabled, an error is written to the <code>Maximo.log</code> file. • If an e-mail record has an error status, that e-mail record is not processed until the error is resolved. • An administrator can use the resubmit function to correct an e-mail record with an error status. When an e-mail is resubmitted, the status is changed to new. The e-mail record is processed as a new e-mail message.
Complete	• Once the workflow process completes successfully, the status of the e-mail record in the staging table is set to complete.

E-mail attachments

The E-mail Listeners application stores attachments from incoming e-mail on the application server. You can view attachments in the **E-mail Processing** tab of the E-mail Listeners application.

The mail server can control the size of the attachment. You can contact your mail server administrator to discuss these controls, and to determine the file types allowed on the E-mail Listeners mail server.

Example

Sally tries to print a file and receives an indecipherable error message. She sends a free form e-mail with a screen capture describing the problem to `help@support.com`, the company site for service desk e-mail requests. The E-mail Listener application retrieves the message from Sally, and creates a service request with identifier 123.

Frank, a service desk agent, reviews service request #123, searches the knowledge base, and finds a solution. He opens the Communications Log containing the initial e-mail from Sally, creates a communication with the solution, and sends it to her.

All details of the interaction between Frank and Sally are stored in the Communications Log for service request #123.

Message thresholds

If the number of messages waiting to be processed exceeds the high message threshold that you set, the application server limits the addition of new messages in the processing queues.

Depending on your message requirements, you might want to type a higher message threshold value. To determine an optimal message threshold setting, you can monitor the messaging in/out queues and the impact of the message threshold setting on system performance. For example, you can lower the threshold value if a higher value is degrading system performance.

If you decide to change the high message threshold setting after the initial configuration, you must open the **Additional Properties** menu in the administrative console. You can then change the threshold value for each child configuration.

Bounced e-mail messages:

Bounced e-mail messages are outbound e-mail messages that cannot be delivered. Large volumes of bounced e-mail messages create excess network traffic and affect the processing of legitimate tickets.

The mail server generates and returns bounced e-mail messages to the e-mail listener account that was specified in the **Send From** field. The e-mail listener treats these messages as service requests.

You can use the following approach to address bounced e-mail messages:

- To preserve the integrity of the primary e-mail listener account, you can create a dedicated e-mail account for bounced e-mail notifications.
- You can base any outbound e-mail notifications on communication templates in which the **Send From** field in the template specifies the dedicated bounced e-mail account. An e-mail message is then generated and sent to the address in the **Send From** field.

Deletion of e-mails from the mail server:

You can manage your e-mail message by specifying a set of rules to mark e-mail messages on the mail server for deletion. You can delete e-mail records with statuses of complete, error, or invalid. When you delete e-mail records in the E-mail Listeners application, the e-mail records are completely removed and cannot be retrieved.

You can use the **E-mail Deleted** option to manage the deletion of e-mails, or delete e-mails manually:

- If you use the **E-mail Deleted** option, the age threshold value and age unit of measure value are used to mark e-mail messages for deletion at set intervals on the mail server. The default value of the age threshold is seven. The default value of the age unit of measure is days.
- If you do not use the **E-mail Deleted** option, read e-mail messages remain on the mail server, regardless of how long messages have been there. In this case, your mail administrator must manage the deletion of e-mail messages from the mail server.

Example

If the age threshold is set to seven and the age unit of measure is days, any read e-mail message that have been on the mail server for seven days are deleted. If you set the age threshold value to zero, any read e-mail marked is deleted from the mail server immediately.

E-mail formats for e-mail listeners

The E-mail Listeners application can process free form e-mail messages or formatted e-mails.

Free form e-mails

Free form e-mails are in plain text and do not follow any specific structure. The E-mail Listeners application first extracts the subject line and body of free form e-mails. Then, it uses the subject line or body to either create a service request or to update an existing ticket.

Free form e-mails are always processed as service requests. If you plan to support other types of tickets, use formatted messages only.

Formatted e-mails

Formatted e-mails use specific structure in the message body to instruct the E-mail Listeners application to manipulate various types of tickets and business objects. Formatted e-mails can be composed using XML tags or text typed in the form of attribute-value pairs. They can perform specific actions, such as changing the status of a business object or querying the business object based on criteria. QUERY, CREATE or UPDATE, and CHANGE STATUS e-mails are all types of formatted e-mails.

To support other business objects, you can create your own workflow process and associated processing logic. A built-in workflow process that supports various types of ticket objects is shipped.

Formatted e-mail keywords:

Formatted e-mails must be carefully composed to ensure that the e-mail listener successfully processes e-mails and that the system performs the necessary actions. You can use specific keywords in the body of the e-mail message to compose a correctly formatted e-mail.

There are two sets of keywords that you can use, depending on whether you want to implement attribute-value pair formatting or XML formatting. The following table specifies the keywords that apply when using attribute-value pairs or XML.

Table 47. Keywords for formatted e-mails

Keyword	Format type	Required	Purpose
#MAXIMO_EMAIL_BEGIN	Attribute-value pairs	Yes	The keyword marks the beginning of formatted e-mail content in an e-mail message.
#MAXIMO_EMAIL_END	Attribute-value pairs	Yes	The keyword marks the end of formatted e-mail content in an e-mail.
LSNRAPPLIESTO	Both	Yes	A value must be provided. The value represents a business object upon which an operation is performed.

Table 47. Keywords for formatted e-mails (continued)

Keyword	Format type	Required	Purpose
LSNRACTION	Both	Yes	A value must be provided. The value can be CREATE, UPDATE, CHANGESTATUS or QUERY. The value specifies the operation to be performed on the business object.
LSNRWHERECONDITION	Both	No	The keyword is used only in e-mails that query a business object. The value represents a valid SQL where condition that is to be applied on the business object.
LSNRRESULTCOLUMNS	Both	No	The keyword is used only in e-mails that query a business object. The value represents one or more columns from the business object. The values are returned in the response e-mail to the query.
&AUTOKEY&	Both	No	The keyword is used as a value for an attribute or an XML tag for e-mail messages that are intended to create business objects. If this keyword is used, the particular attribute is autokeyed using standard system.
&SYSDATE&	Both	No	The keyword is used as a value for an attribute or an XML tag for e-mail messages that are intended to create business objects. If this keyword is used, the value for the particular attribute is a standardized date format as derived from the underlying system database.
<MAXIMOEMAILCONTENT> </MAXIMOEMAILCONTENT>	XML	Yes	This tag is used only in XML formatted e-mails to specify the contents of the e-mail message. This tag serves as the root of the XML document being composed in the e-mail message.

Rules for formatted e-mail messages with attribute-value pairs:

You can use attribute-value pairs to compose formatted e-mail messages.

The following rules apply when you use attribute-value pairs to compose a formatted e-mail message:

- Formatted e-mail messages must contain the #MAXIMO_EMAIL_BEGIN keyword and #MAXIMO_EMAIL_END keyword. If these keywords are not included in the e-mail message, the e-mail message is not processed and an error response e-mail message is sent to the originator of the e-mail message.
- All of the attribute-value pairs specified in the e-mail message must occur together. The e-mail listener ignores any other text typed in the e-mail message. Demarcate the attribute-value pairs in the e-mail message with the #MAXIMO_EMAIL_BEGIN keyword and the #MAXIMO_EMAIL_END keyword. You must place these keywords on a separate line and end with a new line.
- An e-mail message must contain both the #MAXIMO_EMAIL_BEGIN and #MAXIMO_EMAIL_END keywords for the e-mail to be treated as a formatted e-mail. If one or both of the #MAXIMO_EMAIL_BEGIN and #MAXIMO_EMAIL_END keywords are excluded, the e-mail message is treated as free-form text.
- The syntax for the attribute-value pairs is: Field Title#Attribute Name=Value.
- You must place the semi colon character on a new line by itself. This character serves as the separator between one field-value pair and the next field-value pair.
- The field title represents the title for the field as displayed in applications. The attribute name represents the attribute name as specified in the MAXATTRIBUTE table. Typically, this name is the same as the database column name.
- An incoming e-mail message can contain both the field title and the attribute name separated by #, only the field title, or only the attribute name. If only the field title is provided, the e-mail listener attempts to map the field title to the appropriate attribute name before processing the e-mail message. If the e-mail listener cannot map the title or resolve the attribute name, the e-mail message is not processed and an error response e-mail message is sent to the originator of the e-mail message.
- You can place inline attachments before or after the #MAXIMO_EMAIL_BEGIN keyword and #MAXIMO_EMAIL_END keyword.

Rules for formatted e-mail messages with XML:

You can use XML to compose formatted e-mail messages.

The following rules apply when you use XML to compose a formatted e-mail message:

- The XML content of an e-mail message must contain a root element of the form <MAXIMOEMAILCONTENT></MAXIMOEMAILCONTENT>. You must place all other XML tags and values within this root element and are treated as child entities of this root element.
- The syntax for XML is: ATTRIBUTE NAME attribute='Field Title'> ATTRIBUTE NAME attribute='Field Title'>.
- An incoming e-mail message can contain only the attribute name tag. It cannot contain the field title attribute.
- The XML content of an incoming e-mail message is validated only for format. If the XML content is not well formatted, the e-mail entry in the INBOUNDCOMM table is set to an error status and an error notification is sent to both the user and the administrator.

- When processing an XML-formatted e-mail message, the e-mail listener attempts to resolve the attribute name only. If the e-mail listener cannot resolve the attribute name, the e-mail message is not processed and an error response e-mail message is sent to the originator of the e-mail message.
- XML parsers cannot parse the XML content if the standard header is not included in the XML. You must specify XML encoding at the beginning of the body of the e-mail message, before the occurrence of the root element, the MAXIMOEMAILCONTENT tag.
- If XML reserved characters, such as &, occur as a value for any tag in an XML-formatted message, that value must be escaped so that the XML-formatted message constitutes a valid XML. These reserved characters must be escaped using either standard escape sequences or CDATA constructs.
- The keywords &AUTOKEY& and &SYSDATE& must be escaped using standard XML CDATA constructs. For example, <TICKETID><![CDATA[&AUTOKEY&]]></TICKETID>

Examples of formatted e-mail messages using change status function:

You can compose formatted e-mail messages that use the change status function. You compose these messages using attribute-value pairs or XML tags.

Examples

The following table contains examples of formatted e-mail messages that use the change status function. These examples were composed with attribute-value pair and XML tags. You can use these examples as reference or as a template to create formatted e-mail messages in the E-mail Listeners application.

Table 48. Examples of formatted e-mail messages that use the change status function

Description	Example
Attribute-value pair change status function	
Change the status of an existing service request	#MAXIMO_EMAIL_BEGIN LSNRACTION=CHANGESTATUS ; LSNRAPPLIESTO=SR ; CLASS=SR ; TICKETID=SRNUM ; STATUS=INPROG ; #MAXIMO_EMAIL_END
XML change status function	
Change the status of an existing service request	<MAXIMOEMAILCONTENT> <LSNRACTION>CHANGESTATUS</LSNRACTION> <LSNRAPPLIESTO>SR</LSNRAPPLIESTO> <STATUS>QUEUED</STATUS> <TICKETID><SRNUM></TICKETID> <CLASS>SR</CLASS> <SITEID>BEDFORD</SITEID> </MAXIMOEMAILCONTENT>

Examples of formatted e-mail messages using create and update functions:

You can compose formatted e-mail messages that use the create and update functions. You compose these messages using XML tags or attribute-value pairs.

Examples

The following table contains examples of formatted e-mail messages that use the create and update functions. These examples were composed with attribute-value pair and XML tags. You can use these examples as reference or as a template to create formatted e-mail messages in the E-mail Listeners application.

Table 49. Examples of formatted e-mail messages that use the create and update functions

Description	Example
Attribute-value pair create or update function	
Create a service request	#MAXIMO_EMAIL_BEGIN LSNRACTION=CREATE ; LSNRAPPLIESTO=SR ; TICKETID=&AUTOKEY& ; CLASS=SR ; DESCRIPTION= My SR Attribute – value pairs creation TEST ; #MAXIMO_EMAIL_END
Update specific attributes of an existing service request	#MAXIMO_EMAIL_BEGIN LSNRACTION=UPDATE ; LSNRAPPLIESTO=SR ; TICKETID=SRNUM ; CLASS=SR ; DESCRIPTION=Update reported by, priority and classification test. ; REPORTEDPRIORITY=2 ; CLASSSTRUCTUREID=1087 ; #MAXIMO_EMAIL_END
XML create or update function	
Create a service request	<MAXIMOEMAILCONTENT> <LSNRACTION>CREATE</LSNRACTION> <LSNRAPPLIESTO>SR</LSNRAPPLIESTO> <TICKETID><![CDATA[\$AUTOKEY&]]></TICKETID> <CLASS>SR</CLASS> <DESCRIPTION>My XML SR creation e-mail test</DESCRIPTION> </MAXIMOEMAILCONTENT>
Update an existing service request	<MAXIMOEMAILCONTENT> <LSNRACTION>UPDATE</LSNRACTION> <LSNRAPPLIESTO>SR</LSNRAPPLIESTO> <TICKETID>SRNUM</TICKETID> <CLASS>SR</CLASS> <COMMODITYGROUP>IT</COMMODITYGROUP> <COMMODITY>PC</COMMODITY> <DESCRIPTION>My XML update of service group, service, and site field</DESCRIPTION> <SITEID>BEDFORD</SITEID> </MAXIMOEMAILCONTENT>

Examples of formatted e-mail messages using query function:

You can compose formatted e-mail messages that use the query function. You compose these messages using attribute-value pairs or XML tags.

Examples

The following table contains examples of formatted e-mail messages that use the query function. These examples were composed with attribute-value pair and XML. You can use these examples as reference or as a template to create formatted e-mail messages in the E-mail Listeners application.

Table 50. Examples of formatted e-mail messages that use the query function

Description	Example
Attribute-value pair query function	
Query a single record with criteria (LSNRWHERECONDITION keyword)	<pre>#MAXIMO_EMAIL_BEGIN LSNRACTION=QUERY ; LSNRAPPLIESTO=SR ; LSNRRESULTCOLUMNS=TICKETID,DESCRIPTION, REPORTEDBY,COMMODITYGROUP ; LSNRWHERECONDITION=TICKETID='1001' AND SITIED ='BEDFORD' ; #MAXIMO_EMAIL_END</pre>
Query a single record without criteria (LSNRWHERECONDITION keyword)	<pre>#MAXIMO_EMAIL_BEGIN LSNRACTION=QUERY ; LSNRAPPLIESTO=SR ; TICKETID=1002 ; LSNRRESULTCOLUMNS=TICKETID,DESCRIPTION, REPORTEDBY,COMMODITYGROUP ; #MAXIMO_EMAIL_END</pre>
Query multiple records and return selected columns	<pre>#MAXIMO_EMAIL_BEGIN LSNRACTION=QUERY ; LSNRAPPLIESTO=SR ; LSNRRESULTCOLUMNS=TICKETID,DESCRIPTION, REPORTEDBY,INTERNALPRIORITY ,REPORTDATE ; LSNRWHERECONDITION=STATUS ='CLOSED' ; #MAXIMO_EMAIL_END</pre>
Query multiple records and return all columns	<pre>#MAXIMO_EMAIL_BEGIN LSNRACTION=QUERY ; LSNRAPPLIESTO=INCIDENT ; LSNRRESULTCOLUMNS=* ; LSNRWHERECONDITION=REPORTEDBY='LIBERI' ; #MAXIMO_EMAIL_END</pre>
XML query function	

Table 50. Examples of formatted e-mail messages that use the query function (continued)

Description	Example
Query a single record with criteria (LSNRWHERECONDITION tag)	<pre><MAXIMOEMAILCONTENT> <LSNRACTION>QUERY</LSNRACTION> <LSNRAPPLIESTO>PROBLEM</LSNRAPPLIESTO> <LSNRRESULTCOLUMNS>ticketid,description, reportedby,affectedperson,commoditygroup </LSNRRESULTCOLUMNS> <LSNRWHERECONDITION>ticketid in ('1001') </LSNRWHERECONDITION> </MAXIMOEMAILCONTENT></pre>
Query without criteria	<pre><MAXIMOEMAILCONTENT> <LSNRACTION>QUERY</LSNRACTION> <LSNRAPPLIESTO>PROBLEM</LSNRAPPLIESTO> <TICKETID>1003</TICKETID> <LSNRRESULTCOLUMNS>ticketid,description, reportedby,affectedperson,commoditygroup </LSNRRESULTCOLUMNS> </MAXIMOEMAILCONTENT></pre>
Query multiple records and return selected columns	<pre><MAXIMOEMAILCONTENT> <LSNRACTION>QUERY</LSNRACTION> <LSNRAPPLIESTO>PROBLEM</LSNRAPPLIESTO> <LSNRRESULTCOLUMNS>ticketid,description, reportedby,affectedperson,commoditygroup </LSNRRESULTCOLUMNS> <LSNRWHERECONDITION>AFFECTEDPERSON = 'RAMSDALE' AND STATUS = 'QUEUED' </LSNRWHERECONDITION> </MAXIMOEMAILCONTENT></pre>
Query multiple records and return all columns	<pre><MAXIMOEMAILCONTENT> <LSNRACTION>QUERY</LSNRACTION> <LSNRAPPLIESTO>INCIDENT</LSNRAPPLIESTO> <LSNRRESULTCOLUMNS>*</LSNRRESULTCOLUMNS> <LSNRWHERECONDITION>AFFECTEDPERSON = 'SMITH' AND STATUS = 'QUEUED' </LSNRWHERECONDITION> </MAXIMOEMAILCONTENT></pre>

Working with E-mail Listeners

You use e-mail listeners to receive and process incoming e-mail messages. The E-mail Listeners application can monitor multiple e-mail accounts to retrieve messages. It also supports embedded and normal message attachments.

Purging e-mail records from the staging table

As your business needs change, you can purge staging records. If errors are found during the staging process, you can deactivate the listener, purge the staging records, and reactivate the listener to remove the error.

Before you begin

Before you can purge e-mail records from the staging table, you must deactivate the associated e-mail listener definition.

Procedure

1. In the E-mail Listeners definition, select the e-mail listener for which you want to purge e-mail records.
2. Click the **E-mail Processing** tab to view the e-mail records, and from the **Select Action** menu, select **Activate/Deactivate Listener**.
3. Save your changes to deactivate the e-mail listener.
4. From the **Select Action** menu, select **Purge Staging Records**.

5. Click **Yes** to confirm the deletion of the e-mail records.
6. From the **Select Action** menu, select **Activate/Deactivate Listener** to reactivate the e-mail listener.
7. Save your changes.

Example

If you change your mail server implementation or the e-mail address of an e-mail account, you can purge the staging records associated with each account. Alternatively, after initial testing of the application and configurations, you can purge the staging table of any e-mail records before you enable e-mail listener functionality in a production environment.

Related concepts:

“Polling of mail servers for email messages” on page 248

Using email listener definitions, the email Listeners application uses a dedicated cron task to poll the mail server for incoming messages. Cron tasks are Java components that you can set to run automatically and on a fixed schedule. For each email listener definition that you create, a specific instance of the cron task is created and associated with the email listener.

“Staging e-mail messages” on page 249

Using an e-mail listener cron task, the E-mail Listeners application stages e-mail messages to a staging table. When staging an e-mail message, the application saves all information that is required to process the e-mail message and initiate the workflow process. Cron tasks are behind-the-scene jobs that run automatically and on a fixed schedule.

Customizing the e-mail listener preprocessor

You can customize the e-mail listener preprocessor to suit your business needs. The preprocessor determines whether incoming e-mail messages are new requests for help or are updates to existing service requests.

About this task

Java requires that you declare the custom implementation at the beginning of the file. For example: `public class MyPreprocessor implements LSNRPreprocessor`.

Procedure

1. Place the Java class source file in the appropriate Java package where you manage custom Java code.
2. Build your custom Java code into corresponding class files.
3. Build the Enterprise Archive (EAR).
4. Deploy the EAR into the application server to have your code changes take effect.

Related concepts:

“Preprocessors for e-mail listeners” on page 246

A preprocessor determines whether incoming e-mails messages are new requests for help or updates to existing service requests.

Changing the object key delimiter

The object key delimiter value identifies an incoming e-mail as an existing ticket. You can select other characters to represent the object key delimiter other than the default (`##`).

About this task

The ID of the record is called the object key identifier. The object key identifier can be a sequence that the system generates. For example, 1001, 1002, and so on.

Procedure

1. Replace the value with other characters. There are no restrictions. However the delimiter must be unique. Choose infrequently used characters or symbols for delimiters.
2. Place the delimiter before and after the ticket ID (example: SR 1009 is represented as ##1009##).

Example

If + is the delimiter, a user could send an e-mail with the subject line: +1003+ Having problems with printer + network.

The base preprocessor cannot identify the substring because the delimiter symbol occurs multiple times within the subject line. In these circumstances, you must develop your own preprocessor that contains logic to recognize the new delimiter used with e-mails in your business environment.

Related concepts:

“Preprocessors for e-mail listeners” on page 246

A preprocessor determines whether incoming e-mails messages are new requests for help or updates to existing service requests.

Working with e-mail listeners definitions

The E-mail Listeners application uses definitions to poll the mail server for incoming messages and to send e-mail messages in response. The response e-mail messages confirm that a desired operation was performed on behalf of the user.

Creating e-mail listener definitions

When you create an e-mail listener definition, the mail server polls the account that you specify in the e-mail listener definition for incoming e-mail messages. These e-mail messages are processed according to their content.

Procedure

1. In the E-mail Listener application, click **New Listener Definition**.
2. Specify a value for the e-mail address.
3. Optional: Provide a description of the e-mail address.
4. Specify values for the e-mail password, the mail server, the e-mail folder on the mail server that will contain the e-mail messages, and the mail protocol used with the mail server. The port value is provided based on the protocol value. You can change the port value, if necessary.
5. Specify a value for the workflow process to use for the definition.
6. Specify a value for the schedule to set how often you want the server to be polled for incoming e-mail messages. The default frequency is every five minutes. The default values for preprocessor, object key delimiter, cron task name, and cron task instance are provided.
7. Optional: Complete the following steps:
 - To have e-mail messages deleted from the server after they are processed, select the **E-mail Deleted** option.

- Specify a value for the **Age Threshold** field for the length of time that an e-mail message remains on the mail server before being deleted.
 - Specify a value for the age unit of measure.
 - If you configured a Java Messaging Service (JMS) queue to facilitate processing of e-mail messages, select the **Queue Based Processing** check box.
 - If necessary for your configuration, specify a value for the queue connection factory. Specify a value that represents the Java Name and Directory Interface (JNDI) name of the Java component that provides a connection to a queue.
 - If necessary for your configuration, specify a value for the processing queue. Specify a value that represents the name of the queue that is to be used to process e-mail messages for this account.
8. Save the e-mail listener definition.

What to do next

You must activate the e-mail listener definition for incoming e-mail messages to be polled on the mail server. To activate the e-mail listener definition, use the **Activate/Deactivate Listener** action.

Related concepts:

“E-mail listeners definitions” on page 240

As your business needs change, you can change, copy, or deactivate your e-mail listener definitions.

Deleting e-mail listener definitions

You can delete e-mail listener definitions to help simplify e-mail management.

Before you begin

Before you delete an e-mail listener definition, purge the e-mail records from the staging table for that account.

About this task

You can delete e-mail listener definitions with a status of incomplete, invalid, or error. When you delete an e-mail listener definition, the cron task instance associated with that definition is also deleted.

Procedure

1. In the E-mail Listeners application, select the e-mail listener definition that you want to delete.
2. From the **Select Action** menu, select **Activate/Deactivate Listener**.
3. Save the configuration. The e-mail listener definition is deactivated.
4. From the **Select Action** menu, select **Delete Listener**.
5. Save the configuration. The e-mail listener definition and the associated cron task instance are deleted.

Related concepts:

“E-mail listeners definitions” on page 240

As your business needs change, you can change, copy, or deactivate your e-mail listener definitions.

Configuring the queues for WebSphere Application Server

Associating a queue with an e-mail listener can speed up the sequential processing of incoming e-mails, especially when they are arriving in high volume. You can configure the JMS queues for IBM WebSphere Application Server, and associate the queues with listeners.

Procedure

1. Start the WebSphere Application Server.
2. Launch Microsoft Internet Explorer and open the WebSphere administrative console by typing the following URL:
`http://<machine_name>:<port_number>/ibm/console`
For example, type a URL like the following URL:
`http://localhost:9060/ibm/console`
3. At the Welcome screen, enter your User ID then click **Log in**.
4. If necessary, create the MAXIMOSERVER application server:
 - a. In the navigation pane, first, click **Servers**, then **Application Servers**, and then select **New**. Ensure that the application node is ctgNode01.
 - b. In the **Server name** field, enter MAXIMOSERVER, and then click **Next**.
 - c. Click **Next** to accept the default server template.
 - d. Verify that the **Generate Unique Ports** check box is selected and click **Next**.
 - e. From Confirm new server, click **Finish**.
 - f. Save the changes to the master configuration.
5. Change the MAXIMOSERVER JVM heap size properties:
 - a. In the navigation pane, click **Servers** and then select **Application Servers**.
 - b. Click **MAXIMOSERVER**.
 - c. In the Server Infrastructure section, click **Process Definition** in the Java and Process Management section.
 - d. In Additional Properties, click **Java Virtual Machine**.
 - e. Set the **Initial Heap Size** to 1536 for a 32 Bit JVM and to 4096 for a 64 Bit JVM.
 - f. Set the **Maximum Heap** size to 1536 for a 32 Bit JVM and to 4096 for a 64 Bit JVM.
 - g. Click **OK** and then click **Save**.
6. To start the MAXIMOSERVER, click **Servers** and then click **Application Servers**. Select **MAXIMOSERVER** and click **Start**.
7. Click **Service Integration** and then click **Buses**.
8. In the Buses window, click **New**.
9. To add a new service integration bus:
 - a. Type a text description of the new bus in the **Name** field. For example, lsnrjmsbus.
 - b. Clear the **Secure** check box. If you leave this box checked, lsnrjmsbus inherits the Global Security setting of the cell.

- c. In the **High message threshold** field, change the value to a minimum value of 500,000 messages.
- d. Accept all other default settings.
10. Click **Next** and then click **Finish**.
11. Click **Save** to extend the JMS bus setup to the cluster configuration.
12. Confirm that the build completed screen displays the following message:

Bus name, for example, `lsnrjmsbus`.
 Auto-generated, unique ID (UUID), for example, `4BCAC78E15820FED`.
 The Secure field is unchecked.
 High Message Threshold field has a minimum value of 500,000.

Related concepts:

“Queues” on page 248

A queue is an application server component that can facilitate parallel processing. Associating a queue with an e-mail listener can speed up the sequential processing of incoming e-mails, especially when they are arriving in high volume.

“Java Message Driven Bean” on page 248

To enable an e-mail listener to use JMS queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation.

Adding servers to the Java Messaging Service bus for e-mail listeners

You can use IBM WebSphere Application Server to add servers to the Java Messaging Service (JMS) bus for e-mail listeners. Adding servers to the JMS bus is part of configuring JMS queues for e-mail listeners. You configure JMS queues to manage high volumes of e-mail messages that must be processed quickly.

Procedure

1. From the WebSphere administrative console, click **Service Integration** and click **Buses**.
2. In the Buses window, click **lsnrjmsbus** to open the buses.
3. Under the Topology section in the `lsnrjmsbus` window, click **Bus members**.
4. In the Bus members window, click **Add**.
5. In the Add a new bus member window, select the server name **ctgNode01:MAXIMOSERVER** to add to the bus, and then click **Next**.
6. Select **File store**, and then click **Next**.
7. In the Provide the message store properties panel, click **Next**.
8. Click **Finish** and then click **Save**.
9. Select **lsnrjmsbus** and in the **High message threshold** field, change the value to a minimum value of 500,000 messages and then click **Apply**.
10. Select **Synchronize changes with Nodes** and click **Save**.

What to do next

After you add servers to the JMS buses, you create a JMS bus destination for the listener inbound queue.

Related concepts:

“Queues” on page 248

A queue is an application server component that can facilitate parallel processing. Associating a queue with an e-mail listener can speed up the sequential processing of incoming e-mails, especially when they are arriving in high volume.

“Java Message Driven Bean” on page 248

To enable an e-mail listener to use JMS queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation.

Creating the Java Messaging Service bus destination for the listener inbound queue

You need to add a logical address for the listener inbound bus destination queue, `lsnrqin`, within the Java Messaging Service (JMS) bus.

Procedure

1. From the administrative console, click **Service Integration**, and then click **Buses**.
2. In the Buses window, click **lsnrjmsbus**.
3. In the Destination resources section of the `lsnrjmsbus` window, click **Destinations**.
4. In the Destinations window, click **New**.
5. In the Create new destination window, verify that the destination type is queue, and click **Next**.
6. In the Create new queue window, enter `lsnrqin` in the **Identifier** field and Listener Queue Inbound in the **Description** field, then click **Next**.
7. In the Create a new queue for point-to-point messaging window, select **Node=ctgNode01:Server=MAXIMOSERVER** as the bus member to store and process messages for the `lsnrqin` bus destination queue and then click **Next**.
8. In the Confirm queue creation window, click **Finish** to complete the creation of the `lsnrqin` bus destination queue.
9. Select **Buses > lsnrjmsbus > Destinations > lsnrqin**.
10. In the Configuration window, perform the following steps:
 - a. Change the **Maximum failed deliveries** value to 1.
This value is the maximum number of times that you want the system to process a failed messaging attempt before forwarding the message to the exception destination.
 - b. Click **None** as the exception destination value.
11. Click **Apply**, then click **Save**.
12. Select **Synchronize changes with Nodes** and click **Save**.

What to do next

After you add JMS bus destinations for inbound queues for e-mail listeners, you can create the JMS connection factory.

Related concepts:

“Queues” on page 248

A queue is an application server component that can facilitate parallel processing. Associating a queue with an e-mail listener can speed up the sequential processing of incoming e-mails, especially when they are arriving in high volume.

“Java Message Driven Bean” on page 248

To enable an e-mail listener to use JMS queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation.

Creating the Java Messaging Service connection factory

You must add a connection factory to create connections to the associated Java Messaging Service (JMS) provider of point-to-point messaging queues.

Procedure

1. From the WebSphere administrative console, click **Resources**, then click **JMS**, and finally select **Queue Connection Factory**.
2. From the Scope list, select **Cell=ctgCell01** and click **New**.
3. Verify that the **Default Messaging Provider** is selected and click **OK**.
4. Enter the following information:
Name field: lsnrconnfact
JNDI name field: jms/maximo/lsnr/lsnrcf
Bus name field: lsnrjmsbus
5. Click **OK** and then click **Save**.
6. Select **Synchronize changes with Nodes** and click **Save**.

Related concepts:

“Queues” on page 248

A queue is an application server component that can facilitate parallel processing. Associating a queue with an e-mail listener can speed up the sequential processing of incoming e-mails, especially when they are arriving in high volume.

“Java Message Driven Bean” on page 248

To enable an e-mail listener to use JMS queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation.

Creating the listener inbound Java Messaging Service queue

You must create a Java Messaging Service (JMS) queue, lsnrqueue, as the destination for listener inbound point-to-point messages.

Procedure

1. From the administrative console, click **Resources**, then click **JMS**, and finally select **Queues**.
2. From the Scope list, select **Cell=ctgCell01** and then click **New**.
3. Verify the **Default Messaging Provider** is selected and click **OK**.
4. Enter the following information:
Name: lsnrqueue
JNDI name field: jms/maximo/int/lsnr/qin
Bus name field: lsnrjmsbus
Queue name field: lsnrqin

5. Click **OK** and then click **Save**.
6. Select **Synchronize changes with Nodes** and click **Save**.

What to do next

After you create an inbound JMS queue for the e-mail listener, you can activate the queue.

Related concepts:

“Queues” on page 248

A queue is an application server component that can facilitate parallel processing. Associating a queue with an e-mail listener can speed up the sequential processing of incoming e-mails, especially when they are arriving in high volume.

“Java Message Driven Bean” on page 248

To enable an e-mail listener to use JMS queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation.

Activating the listener inbound queue

You must activate the listener inbound queue, **lsnrqueue**, before the queue can receive messages. Activating inbound Java Messaging Service (JMS) queues is part of configuring JMS queues for e-mail listeners.

Procedure

1. From the WebSphere administrative console, click **Resources**, then click **JMS**, and finally select **Activation Specifications**.
2. From the Scope list, select **Cell=ctgCell01** and then click **New**.
3. Complete the General Properties section for the new JMS activation and then click **OK**.
4. Enter the following information:
 - Name** field: **lsnrjmsact**
 - JNDI name** field: **lsnrjmsact**
 - Destination type** field: **Queue**
 - Destination JNDI name** field: **jms/maximo/lsnr/qin**
 - Bus name** field: **lsnrjmsbus**
 - Maximum concurrent endpoints** field: **5**
5. Click **OK** and then click **Save**.
6. Select **Synchronize changes with Nodes** and click **Save**.
7. Stop all IBM-related processes and daemons and then restart these processes for the update to take effect.

Related concepts:

“Queues” on page 248

A queue is an application server component that can facilitate parallel processing. Associating a queue with an e-mail listener can speed up the sequential processing of incoming e-mails, especially when they are arriving in high volume.

“Java Message Driven Bean” on page 248

To enable an e-mail listener to use JMS queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation.

Configuring the Message Driven Bean in WebSphere Application Server

To enable an e-mail listener to use Java Messaging Service (JMS) queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation. Configuring the Message Driven Bean involves removing comment lines from specific sections within the deployment descriptor files of the system.

Before you begin

To complete this task, you need access to the following two files: `ejbjar.xml` deployment descriptor file in the `applications/maximo/mboejb/ejbmodule/META-INF` file path and `ibm-ejb-jar-bnd.xmi` in the `/applications/maximo/mboejb/ejbmodule/META-INF` file path.

Procedure

1. In your installation folder, locate the deployment descriptor file called `ejbjar.xml` under the file path `applications/maximo/mboejb/ejbmodule/META-INF`.
2. Open the file in a text editor and make the following changes:
 - a. Locate the following section and remove the comment lines (`<!--` and `-->`):

```
Email Listener JMS is not deployed by default
<message-driven id="MessageDriven_LSNRMessageBean">
<ejb-class>psdi.common.emailstner.LSNRMessageBEAN</ejb-class>
<transaction-type>Container</transaction-type>
<message-destination-type>javax.jms.Queue</message-destination-type>
</message-driven>-->
```
 - b. Locate the following section and remove the comment lines (`<!--` and `-->`):

```
Email Listener JMS is not deployed by default
<container-transaction>
<method>
<ejb-name>LSNRMessageBean</ejb-name>
<method-name>*</method-name>
</method>
<trans-attribute>Required</trans-attribute>
</container-transaction>-->
```
3. Save the changes that you made to the descriptor file.
4. Locate the file called `ibm-ejb-jar-bnd.xmi` under the file path `/applications/maximo/mboejb/ejbmodule/META-INF` folder.
5. Open the file in a text editor, locate the following section, and remove the comment lines (`<!--` and `-->`):

```
<!-- Email Listener JMS is not deployed by default
<ejbBindings xmi:type="ejbbnd:MessageDrivenBeanBinding"
xmi:id="MessageDrivenBeanBinding_2"
```

```
activationSpecJndiName="lsnrjmsact">
<enterpriseBean xmi:type="ejb:MessageDriven"
href="META-INF/ejbjar.xml#MessageDriven_LSNRMessageBean"/>
</ejbBindings>-->
```

6. Save the changes that you made to the file, then rebuild and redeploy the Enterprise Application Archive (EAR).

What to do next

After you configure the MDBs for e-mail listeners, you can activate workflow processes for e-mail listeners.

Related concepts:

“Queues” on page 248

A queue is an application server component that can facilitate parallel processing. Associating a queue with an e-mail listener can speed up the sequential processing of incoming e-mails, especially when they are arriving in high volume.

“Java Message Driven Bean” on page 248

To enable an e-mail listener to use JMS queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation.

Configuring the Java Messaging Service queues for WebLogic Server

You can use the WebLogic Server to configure Java Messaging Service (JMS) queues for email listeners. You configure JMS queues to manage high volumes of email messages that must be processed quickly.

Procedure

1. Start the WebLogic Server.
2. In the Domain structure, create a file store:
 - a. Expand **Services**, and click **Persistent Stores > New > Create FileStore**.
 - b. In the **Name** field, specify **lsnrstore**, and accept the default value of **AdminServer** for the target.
 - c. In the **Directory** field, specify a folder on the application server computer from which the application server can manage in the store and click **Finish**.
The WebLogic Server can perform read and write operations into the store in this folder. For example, a value for a Windows environment can be **c:\tmp**.
A confirmation message states that the file store was created successfully.
3. In the Domain structure, create a Java Messaging Service server:
 - a. Expand **Services**, and then expand the **Messaging** entry.
 - b. Select **JMS Servers > New** to create a JMS server.
 - c. Name the server **lsnrserver** and select **lsnrstore** as the persistent store.
 - d. In the **Target** field, select **AdminServer** and then click **Finish**.
A confirmation message states that the JMS server was created successfully.
4. In the Domain structure, create a Java Messaging Service module:
 - a. Click **JMS Modules > New**.
 - b. In the **Name** field, specify **lsnrjmsmodule** and click **Next**. Leave the **Descriptor File Name** field and **In Domain** field blank. The application server assigns default values.

- c. Select the **Admin Server** check box and click **Next**.
- d. Select the **Would you like to add resource to this JMS system module** check box and click **Finish**.
A confirmation message states that the JMS module was created successfully.
5. Create a connection factory:
 - a. On the **Configurations** tab on the settings for lsnrjmsmodule page, click **New** in the Summary of Resources table.
 - b. Select **Connection Factory** and click **Next**.
 - c. Enter the following information and click **Next**:
Name field: lsnrconnfact
JNDI name field: jms/maximo/lsnr/lsnrcf
 - d. Verify that the **Targets** field has the following value: AdminServer as selected and click **Finish**.
A confirmation message states that the connection factory was created successfully.
6. Create a Java Messaging Service queue:
 - a. On the **Configurations** tab on the Settings for lsnrjmsmodule page, click **New** in the Summary of Resources table.
 - b. Select **Queue** and click **Next**.
 - c. Enter the following information:
Name field: lsnrqueue
JNDI name field: jms/maximo/int/lsnr/qin
 - d. In the **Template** field, accept the default value of None and then click **Next**.
 - e. In **Targets**, select **lsnrserver** and then click **Finish**.
A confirmation message states that the JMS queue was created successfully.
7. Configure the Java Messaging Service queue:
 - a. On the **Configurations** tab on the Settings for lsnrjmsmodule page, click **lsnrconnfact resource**.
 - b. On the **Configurations** tab on the Settings for lsnrconnfact page, click the **Transactions** tab.
 - c. Select the **XA Connection Factory Enabled** option and then click **Save**.
 - d. In the Change Center, click **Activate Changes**.

Related concepts:

“Queues” on page 248

A queue is an application server component that can facilitate parallel processing. Associating a queue with an e-mail listener can speed up the sequential processing of incoming e-mails, especially when they are arriving in high volume.

“Java Message Driven Bean” on page 248

To enable an e-mail listener to use JMS queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation.

Adding file stores for e-mail listeners - WebLogic Server

You configure Java Messaging Service (JMS) queues to manage high volumes of e-mail messages that must be processed quickly. Adding file stores is part of configuring JMS queues for e-mail listeners.

Procedure

1. In WebLogic Server, create a file store.
2. Specify `lsnrstore` as the name.
3. Accept the default target value of **AdminServer**.
4. Specify a folder on the application server computer from which the store can be managed. The WebLogic Server can perform read and write operations into the store in this folder. For example, a value for an environment can be `c:\tmp`.
5. Click **Finish**.

What to do next

After you add file stores, you add JMS servers for e-mail listeners.

Related concepts:

“Queues” on page 248

A queue is an application server component that can facilitate parallel processing. Associating a queue with an e-mail listener can speed up the sequential processing of incoming e-mails, especially when they are arriving in high volume.

“Java Message Driven Bean” on page 248

To enable an e-mail listener to use JMS queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation.

Adding Java Messaging Service servers for e-mail listeners - WebLogic Server

You configure Java Messaging Service (JMS) queues to manage high volumes of e-mail messages that must be processed quickly. Adding a JMS server is part of configuring JMS queues for e-mail listeners.

Procedure

1. In WebLogic Server, specify `lsnrserver` as the name of the server.
2. Specify the server properties.
3. Create a server, and select **AdminServer** for the target value.
4. Click **Finish**.

What to do next

After you add a JMS server, you add a JMS module for e-mail listeners.

Related concepts:

“Queues” on page 248

A queue is an application server component that can facilitate parallel processing. Associating a queue with an e-mail listener can speed up the sequential processing of incoming e-mails, especially when they are arriving in high volume.

“Java Message Driven Bean” on page 248

To enable an e-mail listener to use JMS queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation.

Adding Java Messaging Service modules for e-mail listeners - WebLogic Server

You configure Java Messaging Service (JMS) queues to manage high volumes of e-mail messages that must be processed quickly. Adding a JMS module is part of configuring JMS queues for e-mail listeners.

Procedure

1. In WebLogic Server, add a JMS module.
2. Specify `lsnrjmsmodule` as the name of the module.
3. Do not specify values for the descriptor file name or in domain. The application server assigns default values.
4. Select the option for the admin server.
5. Select the option to add a resource to the JMS system module.

What to do next

After you add a JMS module, you add a JMS connection factory for e-mail listeners.

Related concepts:

“Queues” on page 248

A queue is an application server component that can facilitate parallel processing. Associating a queue with an e-mail listener can speed up the sequential processing of incoming e-mails, especially when they are arriving in high volume.

“Java Message Driven Bean” on page 248

To enable an e-mail listener to use JMS queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation.

Adding Java Messaging Service connection factories for e-mail listeners - WebLogic Server

You configure Java Messaging Service (JMS) queues to manage high volumes of e-mail messages that must be processed quickly. Adding a JMS connection factory is part of configuring JMS queues for e-mail listeners.

Procedure

1. In WebLogic Server server, create a connection factory.
2. Specify `lsnrconnfact` as the name.
3. Specify `jms/maximo/lsnr/lsnrcf` as the Java Naming and Directory Interface (JNDI) name.
4. Verify that the target value is **AdminServer**, and add the connection factory.

Related concepts:

“Queues” on page 248

A queue is an application server component that can facilitate parallel processing. Associating a queue with an e-mail listener can speed up the sequential processing of incoming e-mails, especially when they are arriving in high volume.

“Java Message Driven Bean” on page 248

To enable an e-mail listener to use JMS queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation.

Adding Java Messaging Service queues for e-mail listeners - WebLogic Server

You configure Java Messaging Service (JMS) queues to manage high volumes of e-mail messages that must be processed quickly. Activating a JMS connection factory is part of configuring JMS queues for e-mail listeners.

Procedure

In WebLogic Server, specify the following information for the queue:

Option	Description
Name	lsnrqueue
Java Naming and Directory Interface (JNDI) name	jms/maximo/int/lsnr/qin
Template	None
Targets	lsnrserver

What to do next

After you add a JMS queue for e-mail listeners, you activate the JMS connection factory.

Related concepts:

“Queues” on page 248

A queue is an application server component that can facilitate parallel processing. Associating a queue with an e-mail listener can speed up the sequential processing of incoming e-mails, especially when they are arriving in high volume.

“Java Message Driven Bean” on page 248

To enable an e-mail listener to use JMS queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation.

Activating Java Messaging Service connection factories for e-mail listeners - WebLogic Server

You configure Java Messaging Service (JMS) queues to manage high volumes of e-mail messages that must be processed quickly. Adding a JMS queue is part of configuring JMS queues for e-mail listeners.

Procedure

1. In WebLogic Server, enable the XA connection factory
2. Save your changes.

3. In the Change Center, activate your changes.

Configuring the Message Driven Bean in WebLogic Server

To enable an e-mail listener to use JMS queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation. Configuring the Message Driven Bean involves removing comment lines from specific sections within the deployment descriptor files of the system.

Before you begin

To complete this task, you need access to the following files: `ejbjar.xml` in the file path `applications/maximo/mboejb/ejbmodule/META-INF` and `weblogic-ejb-jar-bnd.xmi` file in the file path `applications/maximo/mboejb/ejbmodule/META-INF`.

Procedure

1. In your installation folder, locate the file called `ejb-jar.xml` under the file path `applications/maximo/mboejb/ejbmodule/META-INF`.
2. Open the file in a text editor and make the following changes:
 - a. Locate the following section and remove the comment lines (`<!--` and `-->`):

```
<!--Email Listener JMS is not deployed by default
<message-driven id="MessageDriven_LSNRMessageBean">
<ejb-name>LSNRMessageBean</ejb-name>
<ejb-class>psdi.common.emailstner.LSNRMessageBean</ejb-class>
<transaction-type>Container</transaction-type>
<message-destination-type>javax.jms.Queue</message-destination-type>
</message-driven>
-->
```
 - b. Locate the following section and remove the comment lines (`<!--` and `-->`):

```
<!--Email Listener JMS is not deployed by default
<container-transaction>
<method>
<ejb-name>LSNRMessageBean</ejb-name>
<method-name>*</method-name>
</method>
<trans-attribute>Required</trans-attribute>
</container-transaction>
-->
```
3. Save the changes that you made to the file.
4. Locate the file called `weblogic-ejb-jar-bnd.xmi` under the file path `applications/maximo/mboejb/ejbmodule/META-INF`.
5. Open the file in a text editor and locate the following section:

```
<!--Email Listener JMS is not deployed by default
<weblogic-enterprise-bean>
<ejb-name>LSNRMessageBean</ejb-name>
<message-driven-descriptor>
<destination-jndi-name>jms/mro/lsnr/qin</destination-jndi-name>
<connection-factory-jndi-name>jms/mro/lsnr/lsnrcf</connection-factory-jndi-name>
</message-driven-descriptor>
<transaction-descriptor>
<trans-timeout-seconds>600</trans-timeout-seconds>
</transaction-descriptor>
<jndi-name>LSNRMessageBean</jndi-name>
</weblogic-enterprise-bean>
-->
```

 - a. Remove the comment lines (`<!--` and `-->`).

- b. In the section where you removed the comment lines, change the value of the `<connection-factory-jndi-name>` tab to `jms/mro/lsnr/lsnrcf`.
 - c. Save the changes that you made to the file.
6. Rebuild and redeploy the Enterprise Application Archive (EAR) to complete the configuration.

What to do next

After you configure the MDB for e-mail listeners, you configure an e-mail listener to use a JMS queue.

Related concepts:

“Queues” on page 248

A queue is an application server component that can facilitate parallel processing. Associating a queue with an e-mail listener can speed up the sequential processing of incoming e-mails, especially when they are arriving in high volume.

“Java Message Driven Bean” on page 248

To enable an e-mail listener to use JMS queues, you must configure a Java component called a Message Driven Bean that ships with the system. The Message Driven Bean must be configured through a deployment descriptor file that is part of your installation.

Activating workflow processes for e-mail listeners

You configure Java Messaging Service (JMS) queues to manage high volumes of e-mail messages that must be processed quickly. Activating the e-mail listener workflow process to use an e-mail listener as a JMS queue is part of configuring JMS queues for e-mail listeners.

Procedure

1. In the Workflow Designer application, select the **LSNRBP** workflow process.
2. From the **Select Action** menu, select **Enable Process** to validate and enable the workflow process.
3. From the **Select Action** menu, select **Activate Process**.
4. Select the **IBEP** workflow process.
5. From the **Select Action** menu, select **Enable Process** to validate and enable the workflow process.
6. From the **Select Action** menu, select **Activate Process**.

What to do next

After you activate the workflow process for the e-mail listener, you can configure the e-mail listener to use as a JMS queue.

Related concepts:

“Email listeners components” on page 239

The Email Listeners application works together with a CRON task and a workflow process to provide full functionality.

“Predefined workflow process for e-mail listeners” on page 240

You use workflow processes to create steps to guide records for your business process. The E-mail Listener application has a predefined workflow, Listener Business Process (LSNRBP), that is used with the E-mail Listeners application.

“Security settings for e-mail listeners” on page 241

As you can create, update, query, and change the status of tickets, you can configure security settings for e-mail listeners. Using these settings, you can ensure that only authorized users can execute these functions using e-mail messages.

“Communication templates for e-mail listeners” on page 242

To generate standardized e-mail notifications, the E-mail Listeners application uses communication templates to generate notifications that are sent to users and administrators. The types of notifications include confirmations, validation or processing errors, and system errors.

Configuring e-mail listeners to use Java Messaging Service queues

To use an e-mail listener as a Java Messaging Service (JMS) queue, you must configure the e-mail listener. You configure JMS queues to manage high volumes of e-mail messages that must be processed quickly.

Procedure

1. In the E-mail Listeners application, create an e-mail listener definition.
2. On the **Listener** tab, provide the required information, including e-mail address, e-mail password, mail server, protocol, e-mail folder, preprocessor, object key delimiter, workflow process, and schedule.
3. In the E-mail Processing Frequency section, select the queue-based processing option.
4. Specify `jms/maximo/lsnr/lsnrqcf` as the value for the queue connection factory.
5. Specify `jms/maximo/lsnr/qin` as the value for the processing queue.
6. From the **Select Actions** menu, select **Security Settings**.
7. In the Select Security Settings window, click **New Row** and specify the following information:
 - a. Specify **SR** as the value for the Maximo business object.
 - b. Specify **CREATESR** as the value for the Maximo application.
8. Click **OK**.
9. From the **Select Actions** menu, select **Activate/Deactivate Listener**.

Related concepts:

“Email listeners components” on page 239

The Email Listeners application works together with a CRON task and a workflow process to provide full functionality.

“Predefined workflow process for e-mail listeners” on page 240

You use workflow processes to create steps to guide records for your business process. The E-mail Listener application has a predefined workflow, Listener Business Process (LSNRBP), that is used with the E-mail Listeners application.

“Security settings for e-mail listeners” on page 241

As you can create, update, query, and change the status of tickets, you can configure security settings for e-mail listeners. Using these settings, you can ensure that only authorized users can execute these functions using e-mail messages.

“Communication templates for e-mail listeners” on page 242

To generate standardized e-mail notifications, the E-mail Listeners application uses communication templates to generate notifications that are sent to users and administrators. The types of notifications include confirmations, validation or processing errors, and system errors.

Creating communications for e-mail messages

You can create communications that pertain to e-mail records. The communication that you create is sent as an e-mail notification to the recipients that you choose.

Procedure

1. On the **E-mail Processing** tab in the E-mail Listeners application, select the e-mail record for which you want to create a communication.
2. Click **Create Communication** to open the Create Communication window. Based on the e-mail record that you selected, the default values for the recipient, the subject, and the message are provided. You can change these values.
3. Add an attached file.
4. Send the communication to the recipient that you specified.

Related concepts:

“E-mail messages” on page 248

In the E-Mail Listeners application, you can manage your e-mail messages.

“Status of e-mail records” on page 250

Every e-mail record in the staging table has an associated status. These statuses reflect the sequence of actions that the e-mail listener performs on an e-mail record in the staging table.

“E-mail attachments” on page 251

The E-mail Listeners application stores attachments from incoming e-mail on the application server. You can view attachments in the **E-mail Processing** tab of the E-mail Listeners application.

“Bounced e-mail messages” on page 252

Bounced e-mail messages are outbound e-mail messages that cannot be delivered. Large volumes of bounced e-mail messages create excess network traffic and affect the processing of legitimate tickets.

“Deletion of e-mails from the mail server” on page 252

You can manage your e-mail message by specifying a set of rules to mark e-mail messages on the mail server for deletion. You can delete e-mail records with statuses of complete, error, or invalid. When you delete e-mail records in the E-mail Listeners application, the e-mail records are completely removed and cannot be retrieved.

“Message thresholds” on page 251

If the number of messages waiting to be processed exceeds the high message threshold that you set, the application server limits the addition of new messages in the processing queues.

Related reference:

“Examples of formatted e-mail messages using change status function” on page 256

You can compose formatted e-mail messages that use the change status function. You compose these messages using attribute-value pairs or XML tags.

“Examples of formatted e-mail messages using create and update functions” on page 256

You can compose formatted e-mail messages that use the create and update functions. You compose these messages using XML tags or attribute-value pairs.

“Examples of formatted e-mail messages using query function” on page 257

You can compose formatted e-mail messages that use the query function. You compose these messages using attribute-value pairs or XML tags.

Email listener does not process an email

If the email listener does not process an email you can create a DEGUG log statement to isolate the cause of the problem.

Symptoms

A user is expecting a response from an email. The expected email does not arrive in the user's inbox. The user informs the system administrator that the expected email did not arrive.

Causes

- The mail client, such as Microsoft Outlook, is set to a UTF-7 character set encoding. The following error is shown in the log:
java.io.UnsupportedEncodingException: UTF-7 is not a supported format.

- The email contains a digital signature.
- The email sender does not have the authorization to update the service request record which causes the insert and update operations on the business object to fail.
- The email has several embedded messages that are embedded within one another causing the embedded messages to be ignored.

Resolving the problem

1. In the logging application, set the email listener LSNRCRON cron task, to the DEBUG level to produce a detailed log statement.
2. In the System Properties application, add the **mail.debug** property, set it to true, and save the property.
3. Restart the application server for the system property to take effect.
4. After looking at the log statement, set the **mail.debug** property to false.
5. Restart the application server for the system property to take effect.
6. Request the user to resend the original email.

Log statements that support troubleshooting

A DEBUG log level enables the system administrator to isolate the cause of the email listener not processing an email. A log statement contains the following information:

- The number of unread emails
- The subject line of the email
- The header of the email
- The mime type
- The encoding information
- The attachments found in the email
- The attachments that were extracted and saved in the email
- The business object records that were created or updated

Chapter 11. Managing cron tasks

Cron tasks are behind-the-scene jobs that run automatically and on a fixed schedule. You can use the default cron tasks. You can also create your own cron tasks. Cron tasks include activities such as generating preventive maintenance work orders and reordering inventory items on schedules.

Cron task setup overview

You can reschedule cron tasks and change parameter values without stopping and restarting the server. You can create cron tasks and cron task instances, remove cron tasks or cron task instances, and change cron task parameters.

Preexisting cron tasks

There are preexisting cron tasks set up to run scheduled jobs on the system server. Cron tasks are behind-the-scene jobs that run automatically and on a fixed schedule.

The following table describes the preexisting cron tasks. All these cron tasks are set to full access level, except for ESCALATIONS and LSNRCRON (READONLY).

Table 51. Preexisting cron tasks

Cron task	Description
ReorderCronTask	The reorder cron task determines the rules or parameters for scheduled reordering, direct issue, and inventory items.
PMWoGenCronTask	The preventive maintenance work order generation cron task runs and generates scheduled work orders for planned maintenance.
KPICronTask	This cron task generates key performance indicators.
LDAPSYNC	The Lightweight Directory Access Protocol (LDAP) synchronization cron task synchronizes information that is stored in external directory servers for user authentication. The LDAPSYNC cron task supports incremental synchronization.
ESCALATION	The escalation cron task processes escalations to ensure that critical tasks are completed on time.
LSNRCRON	The e-mail listener cron task runs continuously on the application server, and processes inbound e-mail messages through a staging table.
JMSQSEQCONSUMER	This cron task is used by the integration framework to poll the queue.
IFACETABLECONSUMER	This cron task is used by the integration framework to poll the interface tables.

Table 51. Preexisting cron tasks (continued)

Cron task	Description
SwSuiteCronTask	This cron task determines whether the software titles in the Deployed Asset application are present.
ReconciliationCronTask	This cron task runs reconciliation tasks that consist of link and comparison rules. These rules are used to determine how assets are performing relative to the discovered data in the Deployed Asset application. Outputs from this task include a RECONLINK table. This table links assets to their counterpart assets and a reconciliation results table that lists the differences between compared assets and deployed assets.
MeasurePointWoGenCronTask	This cron task generates work orders when meter readings or measurements reach a condition that is defined in the Condition Monitoring application.
FLATFILECONS	This cron task processes inbound flat files.
XMFILECON	This cron task processes inbound XML files.
VMMSYNC	This cron task invokes IBMWebSphere Virtual Member Manager, and invokes WebSphere Virtual Member Manager application programming interfaces (APIs) to populate database tables. The database tables are populated with user group and group membership records. The VMMSYNC cron task supports full synchronization of users and security groups by default, but it can be configured to support the incremental synchronization of users and groups. Before you can use incremental synchronization, you must set the supportChangeLog parameter to native in WebSphere Application Server. If you delete users or security groups in the directory server, the VMMSYNC cron task does not delete the users or groups from the system tables. Use the Security Groups application to delete security groups or the Users application to delete users, and set up an archiving process. Alternatively, you can create a cron task to delete users or security groups.
BBCron	This cron task periodically updates the count for the number of bulletin board messages.

Access levels for cron tasks

Access levels determine what users can do when they are using the Cron Task Setup application.

You can set the access level to one of these options:

- FULL enables users to change or delete the cron task definitions and instances.
- MODIFYONLY enables users only to change the cron task.
- READONLY enables users only to view the cron task.

Cron task parameters

The cron task class file lists parameters. Parameter tables store parameter values for cron task instances.

When you create a cron task instance, the parameter names are retrieved from the cron task class file. For each parameter, a row is added to the parameter table for this instance. When instances are initialized and associated parameters are changed, the instances dynamically obtain the changes from the database.

Example

Table 52. Example of CRONTASKPARAMETER table with data for the ReorderCron cron task

CRONTASKNAME	INSTANCENAME	PARAMETER	VALUE
ReorderCronTask	NA	directisse	
ReorderCronTask	NA	emailto	
ReorderCronTask	NA	ignoreorderpoint	0
ReorderCronTask	NA	leadtime	0
ReorderCronTask	NA	logfile	
ReorderCronTask	NA	storeroom	Nashua
ReorderCronTask	NA	useagreement	1
ReorderCronTask	NA	directissue	
ReorderCronTask	NA	emailto	
ReorderCronTask	NA	ignoreorderpoint	0
ReorderCronTask	NA	leadtime	0
ReorderCronTask	NA	logfile	
ReorderCronTask	NA	storeroom	Central,Bedford
ReorderCronTask	NA	useagreement	1

Instances of cron tasks

An instance is a copy of a cron task that you can change to meet your business needs. You can create multiple instances for each cron task.

You change the attributes or parameters of a cron task instance. Using this method, the same cron task can perform different functions.

Each instance has an entry in the CRONTASKINSTANCE table. The instance includes the following attributes:

- Set schedule string (defines the schedule for this instance)
- Description
- Flag indicating whether the instance is active
- Datetime field indicating the date and time that the load or reload of the cron task is requested (not displayed to users)

- Run as user ID

You have a cron task that performs backups. You can create instances to perform backups at different frequencies, such as daily, weekly, or monthly.

Related tasks:

“Creating cron task instances” on page 285

Before you can run custom cron tasks, you must create a cron task instance. You can create numerous instances for a single cron task, which allows you to fine-tune these tasks to your business needs.

“Copying cron task instances” on page 286

After you create a cron task instance, you can copy it and change it. When you copy a cron task instance, the active status, schedule, run as user ID, and parameter values are also copied.

“Changing cron task instances” on page 286

You can change the attributes or parameter values for a cron task instance to suit your specific needs. You can change the attributes or parameter values without restarting the application server.

“Reloading cron task instances” on page 286

You use the Cron Task Setup application to reload instances of cron tasks.

“Deleting cron task instances” on page 287

In the Cron Task Setup application, you can delete cron task instances that you are no longer using. An instance is a copy of a cron task that you can change to suit your business needs.

Working with cron tasks

Cron tasks are behind-the-scene jobs that run automatically and on a fixed schedule. You can create or change cron tasks.

Creating cron task definitions

A cron task definition consists of a name, class name, access level, and description. You can create cron task definitions to meet the specific needs of your organization.

Before you begin

Create a class file for your custom cron task and package the class file into an Enterprise Application Archive (EAR) file. Then, deploy the EAR file in the application server.

Procedure

1. In the Cron Task Setup application, click **New Cron Task Definition**.
2. Provide a name for the cron task.
3. In the **Class** field, enter the name of the cron task class file. The name is case-sensitive.
4. Specify a value to indicate the level of access that a user has to the cron task.
5. Save the cron task definition.

Deleting cron task definitions

Cron tasks have a definition that includes name, class name, access level, and description. In the Cron Task Setup application, you can delete cron task definitions that you are no longer using.

About this task

The deletion of cron task definitions is based on these rules:

- Cron task definitions with instances or with access levels of READONLY or MODIFYONLY cannot be deleted.
- Cron task definitions must be inactive to be deleted.

Procedure

1. Display the cron task definition that you are deleting.
2. Ensure that the cron task definition has no instances.
3. From the **Select Action** menu, choose **Delete Cron Task**.

Working with instances of cron tasks

You can create multiple instances for each cron task. You can also reload instances of cron tasks, and you can delete instances of cron tasks.

Creating cron task instances

Before you can run custom cron tasks, you must create a cron task instance. You can create numerous instances for a single cron task, which allows you to fine-tune these tasks to your business needs.

Before you begin

The cron task for which you are creating an instance must have a FULL access. You cannot add new instances to cron task definitions that with READONLY access or MODIFYONLY access.

About this task

When you create an instance, the cron task automatically imports a group of associated parameters from the cron task definition. You cannot add new parameters, but you can set parameter values and schedules for each instance.

Procedure

1. In the Cron Task Setup application, select the cron task definition for which you want to create an instance.
2. On the **Cron Task** tab, click **New Row**.
3. Specify a name and a schedule for the cron task instance. The date is shown as a string in the **Schedule** field. Do not change the string in the **Schedule** field. Click **Set Schedule** to change the schedule.
4. Specify a user with the necessary privileges. The user must have access for the actions that the cron task performs.
5. Optional: Select the **Active** check box.
6. Save your changes.
7. From the **Select Action** menu, select **Reload Request**.
8. Select the instance, and click **OK**.

Related concepts:

“Instances of cron tasks” on page 283

An instance is a copy of a cron task that you can change to meet your business needs. You can create multiple instances for each cron task.

Copying cron task instances

After you create a cron task instance, you can copy it and change it. When you copy a cron task instance, the active status, schedule, run as user ID, and parameter values are also copied.

Procedure

1. In the Cron Task Setup application, select the relevant cron task.
2. Click **Duplicate**.
3. Provide a name for the cron task instance.
4. Optional: Change additional information.
5. Save your changes.

Related concepts:

“Instances of cron tasks” on page 283

An instance is a copy of a cron task that you can change to meet your business needs. You can create multiple instances for each cron task.

Changing cron task instances

You can change the attributes or parameter values for a cron task instance to suit your specific needs. You can change the attributes or parameter values without restarting the application server.

Before you begin

Before you change a cron task instance, the instance must have an access level of FULL or MODIFYONLY.

Procedure

1. In the Cron Task Setup application, select the cron task with the instance that you want to change.
2. Click **View Details** for the instance that you want to change.
3. Change applicable information.
4. Save your changes.
5. Determine if the changes that you made require a reload.
6. Optional: If a reload is required, go to the **Select Action** menu and select **Reload Request**.

Example

You change the schedule on which a cron task runs, suspend a task by making it inactive, or change the run as user ID.

Related concepts:

“Instances of cron tasks” on page 283

An instance is a copy of a cron task that you can change to meet your business needs. You can create multiple instances for each cron task.

Reloading cron task instances

You use the Cron Task Setup application to reload instances of cron tasks.

Procedure

1. You must reload an instance only in certain situations. These are the possible situations under which you might need to reload an instance.

Option	Description
Reload Request Not Required	• When switching cron task from Active to Inactive. • Generally, do not submit a reload request after changing parameter values.
Reload Request Required	• When you change any other attributes on the crontask instance, such as schedule and runasuser. • If the parameter value is cached when crontask instance starts or reloads. Check the parameter field to see if this instance occurs.

2. If a reload is required, from the **Select Action** menu, select **Reload Request**.
3. In the Reload Request window, select the instances that you want to reload.
4. Click **OK**.

Related concepts:

"Instances of cron tasks" on page 283

An instance is a copy of a cron task that you can change to meet your business needs. You can create multiple instances for each cron task.

Deleting cron task instances

In the Cron Task Setup application, you can delete cron task instances that you are no longer using. An instance is a copy of a cron task that you can change to suit your business needs.

About this task

The deletion of cron task instances is based on these rules:

- Cron task instances with access levels of READONLY or MODIFYONLY cannot be deleted.
- Cron task instances must be inactive to be deleted.

Procedure

1. Display the instance that you want to delete.
2. Clear the **Active** check box to make the instance inactive.
3. Click **Mark Row for Delete**.

Related concepts:

"Instances of cron tasks" on page 283

An instance is a copy of a cron task that you can change to meet your business needs. You can create multiple instances for each cron task.

Disabling cron tasks on an application server

Active cron tasks run on all application servers in a sequence. In a multi-server environment, you can disable an instance on one or more servers or server clusters. The ReorderCronTask and the PMWoGenCronTask are process-intensive. If the system server is also the corporate print server, you can disable these two cron tasks to reduce the workload of the server.

About this task

You can prohibit all or a selected set of instances from running by modifying the `maximo.properties` file.

Procedure

1. In the `maximo_install\applications\Maximo\properties` directory, open the `maximo.properties` file in a text editor.
2. Navigate to the section Cron Task Manager property. Copy the line `//mxe.crontask.donotrun=crontaskname.instanceName` and paste it in a new line under the text that you copied.
3. In the copied line, remove the comment indicators.
4. Replace the text `crontaskname.instanceName` with the name and instance of the cron task. For example, to disable an instance of the reorder cron task, modify the copied line to say `mxe.crontask.donotrun=ReorderCronTask.ReorderCronTask01`.
5. Save your changes.

What to do next

After you change the `maximo.properties` file, you must rebuild and redeploy the enterprise application archive file.

Viewing hidden cron tasks

READONLY cron tasks are hidden. If necessary, you can view these hidden cron tasks.

Procedure

1. In the Cron Task Setup application, delete **FULL** from the **Access** field.
2. Press **Enter** to view the hidden cron tasks.

Chapter 12. Managing domains

Some fields in the system are associated with select value lists from which users choose an appropriate value. These lists of defined values are known as domains (sometimes referred to as value lists). The system uses many domains in its applications.

Domains overview

You can add domains or modify existing ones to fit with your business practices. After adding domains, additional tasks might be required, depending on the domain and how you want the system to display it.

Applications associated with domains

After you create a domain, you need to apply the domain to an attribute. You use the Classifications, the Database Configuration, and the Application Designer applications to apply and to modify domains after you have created them.

The table that follows lists the applications that are used to apply or to modify domains.

Table 53. Applications used to apply or to modify domains

Application	Description	Example
Classifications	Associate a domain with an attribute in the Attributes table window.	
Database Configuration	Associate a domain with an attribute. Most domains also have a default value, which you would specify in Database Configuration. If the attribute is required, a default value for the domain is also required.	An amount field might be bound to a NUMERIC domain or a status field might be bound to a SYNONYM domain. When you configure the database, the system does not validate the value you insert as the default field value. For example, you can have an Organization called EAGLENA, where the only acceptable domain value is CREW4. You can make the crewid attribute required in the Preventive Maintenance application, give it the default value of CREW2, and configure the database without error. The error, such as CREW2 is not a valid value, appears only when you return to the Preventive Maintenance application to insert a record.

Table 53. Applications used to apply or to modify domains (continued)

Application	Description	Example
Application Designer	Modify the user interface as needed.	If you added an ALN domain for a field, you add the select value button using the application designer application. New Crossover domains might require new fields in the destination application.

Types of domains

Some fields are associated with *Select Value* lists from which you can choose an appropriate value. The lists of defined values are known as *domains* (sometimes referred to as *value lists*).

You can work with the following types of domains:

ALN A list of alphanumeric values.

Crossover

A special table domain in which the system brings back another value (or values) from the specified record.

Numeric

A list of numeric values.

Numeric range

A list of numeric values that you define by specifying a range.

Synonym

These are special, reserved domains in the system. You cannot add or delete synonym domains. You can add new synonym values that are presented to the user.

Table A list of values generated from a table.

Internal and non-internal domains

Some domains are used by the business logic of the system. These domains are called *internal* domains.

If you attempt to edit an internal domain, you encounter the following restrictions:

- You can modify only the description of the internal domain and internal domain values.
- You cannot add or delete a row of values from an internal domain.
- You cannot delete an internal domain.

ALN domains

ALN domains are simple lists of values that use one of the alphanumeric data types.

Example

If your company requires that calendar information is consistent, you can create a list of the days of the week or months of the year. Unlike a SYNONYM domain,

the values in this list are for informational purposes only, the values are not editable.

Crossover domains

Crossover domains return a value from a field in one application to a field in another application. For example, you can return the serial number of an asset in the Assets application to a field in the Item Master application.

When you create a crossover domain, you can specify one or more conditions that must be met before values are returned from a source object. You can define conditions for the source object and the destination object, and if the conditions are met, values are returned to the destination field. You use the Conditional Expression Manager to define conditions for crossover domains.

Order of fields and conditions on crossover domains

If you define conditions on crossover domains in custom applications, carefully consider the order that users scan through the fields. When you define a condition, the condition is evaluated at the moment the value for the attribute is set. If a condition relates to another field with a value that has not been set, values are not returned as expected. For example, this situation might occur due to the order of the layout of the application interface.

Example: Serial number crossover

You want to add a serial number (the destination) to a work order. You want the serial number to cross over from the serial number of the related asset, the source, when the work type is emergency maintenance, the condition. Select the work type before you select the asset number for the crossover to occur as expected.

Numeric range domains

A numeric range domain is a domain that uses one of the numeric data types, but for which you specify a range rather than specific values.

You can specify the following kinds of ranges:

Discrete

A range with a defined interval between values, for example, a range from 0 to 10, with valid values 0, 2, 4, 6, 8, and 10. The interval in this example is 2.

Continuous

A range within which any value that satisfies the data type is valid. For example, in a range of 1 to 6 with a decimal data type and scale of 2, values such as 1, 1.03, 2.14, 3, 4.73, 5.2, and 6 are all valid. The interval in this example is null (no value in the **Interval** field).

You cannot create lookups for numeric range domains. Therefore, consider the types of values a user or automated process normally enter into a field with their type of domain. An invalid value results in an error message.

Numeric ranges with more than one segment

You might use multiple segments in your numeric range domain for various reasons. Here are two examples:

- You want to define a measurement range that is more precise at low measurements than at high measurements and that correspond with meter readings. You can define three segments:
 Segment 1: minimum 0, maximum 0.8, interval 0.2, resulting valid values, 0, 0.2, 0.4, 0.6, 0.8
 Segment 2: minimum 1, maximum 9, interval 1, resulting valid values 1, 2, 3, 4, 5, 6, 7, 8, 9
 Segment 3: minimum 10, maximum 30, interval 5, resulting valid values 10, 15, 20, 25, 30

A user or automated process would always enter one of those values.

- You want values inserted into a field only if the reported values are beyond the accepted normal range. For example, meter readings could be above or below the accepted tolerances. You can define two segments:
 Segment 1: minimum 0, maximum 9.9, interval null. Any reading between 0 and 9.9 could be entered.
 Segment 2: minimum 20.1, maximum null, interval null. Any reading of 20.1 or higher could be entered.

Meter readings from 10 and 20 would not be recorded.

Synonym domains

Synonym domains are special domains that are reserved by the system. You cannot add new synonym domains, but you can add new synonym values that are presented to the user.

An example of a synonym domain is work order status. The system has several values to reflect status: APPR (Approved), CAN (Canceled), CLOSE (Closed), COMP (Completed), WAPPR (Waiting on Approval), and others. Each work order status has an internal value, used by the system in its business rules, and a value that users see and choose from. You cannot add a new internal value. You can add a synonym, the value that is presented to the user.

Example: Synonym values for WAPPR (Waiting on Approval)

Suppose your company procedure requires two people to approve a work order. You could add synonym values for the internal WAPPR value. You could then present two different values to the user, for example, WAPPRMAN and WAPPRVP, to represent approvals at the manager and vice president level.

Related tasks:

“Creating synonyms of internal values” on page 303

You can create synonyms of internal values so that you can present different values to users based on your business needs.

“Deleting synonyms of internal values” on page 304

You can delete a synonym of an internal value when it is no longer useful.

TABLE domains

TABLE domains are dynamic sets of values based on the values of another object.

Example

You can use a TABLE domain to present a valid list of records from the PERSON table to be typed in the **OWNER** field on a record.

Foreign keys and TABLE domains

Using domains, you can create a foreign key from a system level for site-level or for organization-level objects.

Example

If you want to add a new attribute for assets on the TKTEMPLATE object, perform the following steps:

1. Create a TABLE domain.
2. Add the assetnum attribute to the TKTEMPLATE.
3. Add a relationship to the asset table.
4. Add the attribute siteid to TKTEMPLATE.

Domains and organizations or sites

Applying domains to the organization or site level (by typing appropriate values in the organization field and site field) might create unintentional access restrictions. Leaving these fields blank is the default, and stores domains at the system level.

Example one

When you specify an organization or site for domain values, note where the domain is being used.

In the Labor application, you use the SKILLLEVEL domain on the CRAFTSKILL object. You specify both the organization and site values for domain values. When you access the Labor application and lookup the skill level, you do not see your values. This issue occurs because you specified a site, and the object that is using the domain is at the organization level. To fix this problem, remove the site from the domain value.

Example two

You might want to leave the organization field and site field empty for all values (users in all organizations and sites can access them). If you specify an organization or site for one value, you must specify an organization or site for all values (users in the specified organizations or sites can access them).

Otherwise, complicated outcomes can result. For example, you can set an organization with the following domain values:

Table 54. Sample domain setup with specific organizations

Value	Organization
GREEN	A
BLUE	B
RED	(No organization specified)

With the domain configuration above, you get the following results:

- Records in Organization A can access GREEN only.
- Records in Organization B can access BLUE only.
- Records in other organizations (other than Organization A and Organization B) can access RED only.

Working with domains

You use domains to choose values associated with specific fields. The lists of defined values are known as domains (sometimes referred to as value lists). The system uses many domains in its applications.

Adding alphanumeric domains

You add an alphanumeric (ALN) domain when you want to add a domain that uses one of the alphanumeric data types.

About this task

Because you can use a domain with multiple fields, the length you specify must be less than or equal to the length of the shortest field with which you intend to use the domain. For example, if you want to use the domain with three fields of lengths 8, 10, and 12, specify a length of 8 or less. If you specify a length greater than the field the domain is used in, you cannot assign the domain to the attribute in the Database Configuration application. Alternatively, you can specify a greater length in the Domains application and use the Database Configuration application to change the length of the field that uses the domain.

Procedure

1. Open the Domains application.
2. At the bottom of the Domains table window, click **Add New Domain** and select **Add New ALN Domain**.
3. In the **Domain** field, specify a name for the domain.
4. In the **Description** field, type a short description for the domain.
5. In the **Data Type** field, specify a data type for the domain.
6. In the **Length** field, specify a length that is equal to or less than the length of the field that uses the domain. For example, if you are adding a domain for a field in the Assets application that has a length of 12, then specify a length of 12.
7. Click **New Row**.
8. Fill the **Value** and **Description** fields.
9. Optional: Apply a domain value to a specific organization or site by entering the relevant values in the **Organization** and **Site** fields. Domains are applied at the system level by default.

You can leave the **Organization** and/or **Site** fields empty for all values so that users in all organizations and sites can access them. Or you could specify an organization and/or site for all values so that only the users in the specified organizations and/or sites can access them. Otherwise, complicated outcomes can result.

For example, if you enter a value RED with no organization specified, a value GREEN with Organization A specified, and a value BLUE with Organization B specified, then records in Organization A has access to GREEN only. Records that are in Organization B has access to BLUE only. Records that are in other organizations has access to RED only. When you specify an organization and/or site for one value, records in that organization and/or site no longer have access to values that have no organization or site specified.

10. Optional: Click **New Row** again to add more values.
11. Click **OK**.

What to do next

After you add a domain, you might still have several tasks to perform, depending on the domain and how you want to display it.

In the Classifications application, you associate a domain with an attribute in the Attributes table window. No further configuration is needed.

If you use a domain in any other context, adding a domain requires additional tasks:

- Associate the new domain with an attribute in the Database Configuration application.
- Configure the database in the Database Configuration application.
- Use the Application Designer application to modify the user interface as needed. For example, if you add an ALN domain for a field, you must add the select value button. New crossover domains might require new fields in the destination application.

Related concepts:

“Numeric range domains” on page 291

A numeric range domain is a domain that uses one of the numeric data types, but for which you specify a range rather than specific values.

Adding crossover domains

You add a crossover domain when you want to add a domain that returns a value from a field in one application to a field in another application.

About this task

Ensure that the SQL statements that you use in this procedure are valid. The application does not validate SQL statements.

Procedure

1. Open the Domains application.
2. In the Domains table window, click **Add New Domain** and select **Add New CROSSOVER Domain**.
3. In the **Domain** field, specify a name for the domain.
4. In the **Description** field, type a short description for the domain.
5. In the Crossover Domain table window, click **New Row**.
6. Define details for the crossover domain:

Field	Description
Object	Select the name of the object that contains the attribute that you want to create a domain from. For example, to obtain values from the ASSET object, select ASSET.

Field	Description
Validation Where Clause	If the value to be validated by this domain is considered valid, type the part of the clause that when queried against the object in the Object field returns at least one record. For example, if you want a field named Z (attribute Z) to contain values from the assetnum field in the Assets application, you would type: ASSETNUM = :Z (the colon represents the bind variable).
List Where Clause	Type the part of the clause that specifies the value that you want to select based on the validation WHERE clause. For example, to select asset records that begin with the numbers 114, type ASSETNUM LIKE '114%'

7. Optional: Specify a group value and a key value to select an error message to display when domain validation fails. You define values for error messages in the Database Configuration application.
 - a. In the **Error Message Group** field, specify the group value of the error message.
 - b. In the **Error Message Key** field, specify the key value of the error message.
8. Optional: To apply a domain to a specific organization or site, specify values in the **Organization** and **Site** fields. To allow access to users in all organizations and sites, clear all values from these fields.
9. Select fields for the crossover domain:
 - a. In the Crossover Fields table window, click **New Row**.
 - b. In the **Source Field** list, select an attribute of the object that you specified in the **Object** field. This attribute represents the field from which you want to return values.
 - c. In the **Destination Field** text box, specify the field to which you want values to be returned.
 - d. Optional: Select the **Accept NULL Value** check box to copy the value from the source field when the value of the target attribute is empty. This function overwrites the previous value.
 - e. Optional: Select the **No Overwrite** check box if you want to copy the value of the source attribute when it is null.
10. Optional: Specify conditions for the source object or the destination object, or both:
 - a. In the **Condition on Source** field, select a condition or use the Conditional Expression Manager to build the condition for the source object of the crossover. The source object must meet the condition before the crossover occurs.
 - b. In the **Condition on Destination** field, select a condition or use the Conditional Expression Manager to build the condition for the destination object of the crossover. The destination object must meet the condition before the crossover occurs.

- c. In the **Sequence** field, type a numeric value to specify the order that the crossover occurs when multiple crossovers are defined. Crossovers with lower values occur before crossovers with higher values.
11. Click **OK** to add the crossover domain to the database.

What to do next

After you add a domain, you might still have several tasks to perform, depending on the domain and how you want to display it.

In the Classifications application, you associate a domain with an attribute in the Attributes table window, and no further configuration is needed.

If you use a domain in any other context, adding a domain requires additional tasks:

- Associate the new domain with an attribute in the Database Configuration application.
- Configure the database in the Database Configuration application.
- Use the Application Designer application to modify the user interface as needed. For example, if you add an ALN domain for a field, you must add the select value button. New crossover domains might require new fields in the destination application.

Related concepts:

“Crossover domains” on page 291

Crossover domains return a value from a field in one application to a field in another application. For example, you can return the serial number of an asset in the Assets application to a field in the Item Master application.

Adding numeric domains

You add a numeric domain when you want to add a domain that uses one of the numeric data types.

About this task

Because you can use a domain with multiple fields, the length that you specify here must be less than or equal to the length of the shortest field with which you use the domain. For example, if you want to use the domain with three fields of lengths 8, 10, and 12, specify a length of 8 or less for the domain. If you specify a length greater than the field the domain is used in, you are not able to assign the domain to the attribute in the Database Configuration application. Alternatively, you can specify a greater length in the Domains application. You can use the Database Configuration application later to change the length of the field that uses the domain.

Procedure

1. Open the Domains application.
2. At the bottom of the Domains table window, click **Add New Domain** and select **Add New NUMERIC Domain**.
3. In the **Domain** field, specify a name for the domain.
4. In the **Description** field, type a short description for the domain.
5. In the **Data Type** field, specify a data type for the domain.
6. Depending on the data type the **Length** field might or might not be editable. If editable, specify a length that is equal to or less than the length of the field

that uses the domain. For example, if you are adding a domain for a field in the Assets application that has a length of 12, then specify a length of 12.

7. For Decimal data type only, in the **Scale** field, specify a scale value if different from the default, 2.
8. Click **New Row**.
9. Fill the **Value** and **Description** fields.
10. Optional: Apply a domain value to a specific organization or site by specifying values in the **Organization** and **Site** fields. Domains are applied at the system level by default.

You might want to leave the Organization and/or Site fields empty for all values so that users in all organizations and sites can access them. Or you could specify an organization and/or site for all values so that only users in the specified organizations and/or sites can access them. Otherwise, complicated outcomes can result.

For example, if you enter a value 100 with no organization specified, a value 300 with Organization A specified, and a value 500 with Organization B specified, then records in Organization A have access to 300 only. Records that are in Organization B have access to 500 only. Records that are in other organizations have access to 100 only. Once you specify an organization and/or site for one value, records in that organization and/or site no longer has access to values that have no organization/site specified.
11. Optional: Click **New Row** again to add more values.
12. Click **OK**.

What to do next

After you add a domain, you might still have several tasks to perform, depending on the domain and how you want to display it.

In the Classifications application, you associate a domain with an attribute in the Attributes table window, and no further configuration is needed.

If you use a domain in any other context, adding a domain requires additional tasks:

- Associate the new domain with an attribute in the Database Configuration application.
- Configure the database in the Database Configuration application.
- Use the Application Designer application to modify the user interface as needed. For example, if you add an ALN domain for a field, you must add the select value button. New crossover domains might require new fields in the destination application.

Related concepts:

“Numeric range domains” on page 291

A numeric range domain is a domain that uses one of the numeric data types, but for which you specify a range rather than specific values.

Adding numeric range domains

You add a numeric range domain when you want to add a domain that uses one of the numeric data types. However, you specify a range rather than specific values.

About this task

Because you can use a domain with multiple fields, the length that you specify here must be less than or equal to the length of the shortest field you use the domain with. For example, if you want to use the domain with three fields of lengths 8, 10, and 12, specify a length of 8 or less for the domain. If you specify a length greater than the field the domain is used in, you cannot assign the domain to the attribute in the Database Configuration application. Alternatively, you can specify a greater length in the Domains application. You can use the Database Configuration application later to change the length of the field that uses the domain.

Procedure

1. Open the Domains application.
2. At the bottom of the Domains table window, click **Add New Domain** and select **Add New NUMERIC RANGE Domain**.
3. In the **Domain** field, specify a name for the domain.
4. In the **Description** field, type a short description for the domain.
5. In the **Data Type** field, specify a data type for the domain.
6. Depending on the data type, the **Length** field might or might not be editable. If editable, specify a length that is equal to or less than the length of the field that uses the domain. For example, if you are adding a domain for a field in the Assets application that has a length of 12, then specify a length of 12.
7. For DECIMAL data type only, in the **Scale** field, specify a scale value if different from the default, 2.
8. Click **New Row** and fill the following fields. Each row defines a range.

Field	Description
Range Segment	Specify an identifying number for the segment. For example, 1 if the range has just one segment, or 1, 2, or 3 if the domain has three segments. For a range with regular intervals or a simple continuous range, you need only one segment.
Range Minimum	Specify the lowest value in the range, for example, 10 in the range 10 to 50. You can leave either the Range Minimum field or the Range Maximum field empty (null value), but not both.
Range Maximum	Specify the highest value in the range, for example, 50 in the range 10 to 50. You can leave either the Range Minimum field or the Range Maximum empty (null value), but not both.

Field	Description
Interval	For a range with discrete values, specify the value of the interval that separates the values you want to appear in the list. For example, the interval is 10 in the range 10 - 50 if you want the values 10, 20, 30, 40, and 50 to be valid. For a continuous range of values between the minimum and maximum values, leave this field null (empty). Continuous is relative to the data type. For example, if the data type is INTEGER and the range is 1 - 5, then only 1, 2, 3, 4, and 5 are valid values.

9. Optional: Apply a domain value to a specific organization or site by specifying values in the **Organization** and **Site** fields. Domains are applied at the system level by default.
You might want to leave the **Organization** and/or **Site** fields empty for all range segments so that users in all organizations and sites can access them. Or you could specify an organization and/or site for all range segments so that only users in the specified organizations and/or sites can access them.
10. Optional: Click **New Row** again to add ranges for additional segments.
11. Click **OK**.

What to do next

After you add a domain, you might still have several tasks to perform, depending on the domain and how you want to display it.

In the Classifications application, you associate a domain with an attribute in the Attributes table window, and no further configuration is needed.

If you use a domain in any other context, adding a domain requires additional tasks:

- Associate the new domain with an attribute in the Database Configuration application.
- Configure the database in the Database Configuration application.
- Use the Application Designer application to modify the user interface as needed. For example, if you add an ALN domain for a field, you must add the select value button. New crossover domains might require new fields in the destination application.

Related concepts:

“Numeric range domains” on page 291

A numeric range domain is a domain that uses one of the numeric data types, but for which you specify a range rather than specific values.

Adding table domains

You add a table domain when you want to add a domain that draws its values directly from a column in the database. This process creates a dynamic value list because the values it draws from the database might change.

Procedure

1. Open the Domains application.
2. At the bottom of the Domains table window, click **Add New Domain** and select **Add New TABLE Domain**.
3. In the **Domain** field, specify a name for the domain.
4. In the **Description** field, type a short description for the domain.
5. Click **New Row** and fill the following fields.

Field	Description
Object	Specify the name of the object. The object you want is the object containing the attribute from which you want to create a domain. For example, to obtain values from the ASSET object, specify ASSET.
List Where Clause	Type the part of the clause that specifies the values that you want to select based on the validation WHERE clause. For example, to select asset records that begin with the numbers 114, type: assetnum like '114%' Attention: The system does not validate your entry for syntax or any other errors. Be sure that you have typed a correct WHERE clause. If you make errors, errors do not become apparent until you configure the database.
Validation Where Clause	If the value to be validated by this domain is considered valid, type the part of the clause that when queried against the object in the Object field should return at least one record. Usually, the clause involves a bind variable for the field that uses this domain for validation. The bind variable is represented by a colon (:) followed by the field name. For example, if you want a field named Z (attribute Z) to contain values from the assetnum field in the Assets application, type: assetnum = :z

6. Optional: Specify a group value and a key value to select an error message to display when domain validation fails. You define values for error messages in the Database Configuration application.
 - a. In the **Error Message Group** field, specify the group value of the error message.
 - b. In the **Error Message Key** field, specify the key value of the error message.
7. Optional: Apply a domain value to a specific organization or site by specifying values in the **Organization** and **Site** fields. Domains are applied at the system level by default.

You might want to leave the **Organization** and/or **Site** fields empty for all values so that users in all organizations and sites can access them. Or you could specify organization and/or site for all values so that only users in the specified organizations and/or sites can access them.

8. Optional: Click **New Row** again to add more rows.
9. Click **OK**.

What to do next

After you add a domain, you might still have several tasks to perform, depending on the domain and how you want to display it.

In the Classifications application, you associate a domain with an attribute in the Attributes table window, and no further configuration is needed.

If you use a domain in any other context, adding a domain requires additional tasks:

- Associate the new domain with an attribute in the Database Configuration application.
- Configure the database in the Database Configuration application.
- Use the Application Designer application to modify the user interface as needed. For example, if you add an ALN domain for a field, you must add the select value button. New crossover domains might require new fields in the destination application.

Related concepts:

“TABLE domains” on page 292

TABLE domains are dynamic sets of values based on the values of another object.

“Foreign keys and TABLE domains” on page 293

Using domains, you can create a foreign key from a system level for site-level or for organization-level objects.

Associating domain values with conditions

You can associate domain values with a condition to limit the number of values available to users for certain conditions such as statuses, priorities, and assets. This association ensures that only users determined by you can set certain statuses, change the priority of work orders, and so on.

Before you begin

Before you associate a domain value with a condition, you must set up the domain values for the ALN, synonym, and numeric domains. You must then define a condition for when particular domain values should appear in a domain.

About this task

You can create a condition in the condition library of the Conditional Expression Manager application. If you modify the condition for a domain value after the value has been selected for a field, the existing values are not revalidated. There is no current record on the Search or List pages so conditional values are not supported on these pages.

Procedure

1. Open the Domains application and select the **Edit Detail** icon for the domain containing the values you want to associate with a condition. Depending on the domain type selected, an ALN, numeric, synonym, table, or crossover window opens from where you can add or change conditions.
2. Add or change the conditions as necessary.
3. Optional: If you want to add the same conditions for multiple domain values, select **Set Conditions** and select a condition from the condition library. You can apply the same condition to all of the selected values.
4. Optional: If you want domain values to be used by multiple objects, specify the objects in the **Object Name** field.
5. Save the record.

Creating synonyms of internal values

You can create synonyms of internal values so that you can present different values to users based on your business needs.

Procedure

1. Open the Domains application.
2. Find the synonym domain for which you want to add a synonym value, for example WOSTATUS.
3. Click **Edit Detail**. In the Synonym Domain table window, the set of current values: the Internal Value is used by the system and must be unique. The Value and its description are what users see.
4. Click **New Row** and enter values in the fields:

Field	Description
Internal Value	Specify the internal value for which you want a synonym. For example, in the WOSTATUS domain, if you want to create a synonym for WAPPR called WAIT, you would specify the WAPPR internal value.
Value	Specify the synonym value you want a user to see. For example, WAIT.
Description	Type a description of the synonym to differentiate it from the existing value.

5. Optional: In the **Default** field, select the check box if you want the system to use the new synonym value by default. Each internal value can have only one default synonym value.

For example, the system inserts WAPPR (Waiting for Approval) as the status when you create a work order. You add a synonym value, WAIT. If you want the system to insert WAIT instead of WAPPR, then make WAIT the default.

6. Optional: Apply a domain value to a specific organization or site by specifying values in the **Organization** and **Site** fields. Domains are applied at the system level by default.

You might want to leave the **Organization** and/or **Site** fields empty for all values so that users in all organizations and sites can access them. Or you

could specify organization and/or site for all values so that only users in the specified organizations and/or sites can access them.

If you create a synonym value and specify a site or organization, and then click **OK**, the system automatically inserts duplicate rows for the other values. The duplicate rows have the site or organization you specified for the new value that you created. For example, the MRTYPE domain has three values: RECURRING, STANDARD, and TRANSFER (each with an internal value of the same name). If you create a synonym value, REGULAR, with the internal value of STANDARD and specify Organization B, the system automatically creates (when you click **OK**) the additional synonym values RECURRING and TRANSFER with Organization B specified.

Important: After you have implemented the system and inserted records, you generally must not add a synonym value with a site or organization specified. This action can invalidate existing data. If you must add a synonym value with a site or organization specified, you must also add synonym values for all the existing values that are in the database. For example, if you added a synonym value for STANDARD with an organization or site specified, you must also add synonym values for RECURRING and TRANSFER (if they have been used on records) with the same organization or site specified.

7. Click **New Row** to add more synonyms.
8. Click **OK**.

Related concepts:

“Synonym domains” on page 292

Synonym domains are special domains that are reserved by the system. You cannot add new synonym domains, but you can add new synonym values that are presented to the user.

Deleting synonyms of internal values

You can delete a synonym of an internal value when it is no longer useful.

About this task

In synonym domains, you can delete synonym values, with the following conditions:

- You can delete only non-default synonym values.
- If you delete one synonym value for which you have specified an organization or site, you must delete all other values for which you have specified that organization or site.
- Do not delete synonym values if you have records referencing them.

Procedure

1. Open the Domains application.
2. In the Domains table window, find the domain with the synonym value that you want to delete and click **Edit Detail**.
3. In the SYNONYM Domain table window, click the **Mark Row for Delete** icon. You can mark multiple rows.
To cancel a deletion, click the **Undo Delete** icon.
4. Click **OK**.

Related concepts:

“Synonym domains” on page 292

Synonym domains are special domains that are reserved by the system. You cannot add new synonym domains, but you can add new synonym values that are presented to the user.

Deleting domains

You can delete most domains when they are no longer useful.

About this task

You cannot delete a synonym domain.

You can delete other types of domains, but not if the domain is assigned to an attribute. To delete the domain, you must first disassociate the attribute.

Deleting a domain does not affect values that have already been inserted on records. For example, if a user inserts a value in a field using a domain select value list, that value remains on the record. The record remains on the record even if the domain is later disassociated from the attribute and deleted.

Procedure

1. Open the Domains application.
2. In the Domains table window, find the domain that you want to delete and click the **Mark Row for Delete** icon.
3. In the warning message, click **Yes**.
4. Optional: You can click the **Undo Delete** icon to cancel the deletion.
5. Click **Save Domain**.

Chapter 13. Configuring and administering attached documents

You can attach relevant information to a record (or to a task on a record) in the form of a file or a URL address. Attached documents can be located on your company network, on the Internet, or in a Document Management System. Documents can include text files, spreadsheets, and diagrams.

About this task

When you configure the **Attached Documents** action, you integrate the location of a stored document file with the location that you specified in the system. You can configure the system to store attached document files on the same server as the application server that is running the system. You can also store attached documents on other servers.

If you have a document management system, you can integrate it with the **Attached Documents** action. Integrating a document management system requires code changes and programming skills.

Attached Documents overview

You use the **Attached Documents** action, found in most system applications, to create a document library and to organize documents into folders. The system includes default folders. You can also create more folders or organize the folders into functional categories.

The Maximo database includes the following folders that you can use for this purpose.

Table 55. Folders in Maximo where you can organize your documents

Folder	Contents
Attachments	Text files
Diagrams	Flow charts or part diagrams
Images	Graphic images, such as pictures of assets

Administrators maintain the library, create folders as needed, and specify the folders available for each application. Additional folders might include permits, part sheets, photographs, procedures, drawings. You can attach a document to a record even when the document is outside the document library.

Configuration of attached documents

The system uses WebSphere Application Server or WebLogic Server. Using one of these application servers, you can configure the Attached Documents action on a single computer or on multiple computers.

The computers from which you access attached documents must have the relevant applications installed on them. For example, to view a Microsoft Word document, a workstation must have Word installed on it.

Configuration of attached documents for a single computer

For WebSphere Application Server and WebLogic Server, the single computer has a specific configuration and specifications. This configuration assumes that your computer is running a Windows or UNIX scenario.

For this configuration, the application server and the HTTP server are on the same computer. You also must store document files on the same single computer on which the application server and HTTP server are running. The configuration shown in the following figure stores attached document files on the same system as the application server that runs the system:

The following graphic depicts a single computer configuration for attached documents.




Figure 6. Single computer configuration for attached documents

Configuration of attached documents for two computers and a local Hypertext Transfer Protocol server

When you use two computers, a local Hypertext Transfer Protocol (HTTP) server, and WebSphere Application Server on either Windows or UNIX, certain configuration specifications apply.

The following table lists the configuration specifications that apply for two computers, a local HTTP server, on Windows or UNIX.

Table 56. Two computers, local HTTP server, on Windows or UNIX

Configuration	Specifications
- You store document files on a different computer than the application server that runs the system. - The HTTP server is on the application server. - You map a drive on the application server to point to the drive on the document file server (Windows only). - You mount the network file system that contains the document files from the document file server computer onto the application server (UNIX only).	For Windows: - H is a mapped drive on the application server computer that runs the system. - D is a drive on the computer that stores the documents. - Drive letters, and file names and directory names are case-sensitive. For UNIX: /d01 is the NFS mount point on the application server for the file system /home on the document storage computer. - File names and directory names are case-sensitive.

The following figure shows a two-computer configuration with a local HTTP server.




Figure 7. Two-computer configuration with a local HTTP server

Configuration of attached documents for two computers and a dedicated Hypertext Transfer Protocol server

When you use two computers with a dedicated Hypertext Transfer Protocol (HTTP) server on either Windows or UNIX, the following configuration and specifications apply. This configuration applies to the WebSphere Application Server platform or to the WebLogic Server platform.

The following table lists the configuration specifications that apply to two computers, with a dedicated HTTP server, on Windows or UNIX

Table 57. Two computers, dedicated HTTP server, on Windows or UNIX

Configuration	Specifications
- You store document files on a different computer than the application server that runs the system. - The HTTP server (such as Apache or Microsoft Internet Information Services) is on the computer that stores the document files. - You map a drive on the application server to point to the drive on the Document File/HTTP server (Windows only). - You mount the network file system that contains the document files from the document file server onto the application server (UNIX only).	For Windows: - H is a mapped drive on the application server computer that runs the system. - D is a drive on the computer that stores the documents, and runs an HTTP server. - Drive letter, file names, and directory names are case-sensitive. For UNIX: /d01 is an NFS mount point on the application server for the file system /home on the HTTP server. - File names and directory names are case-sensitive.

The following graphic depicts the configuration for two computers with a dedicated document file/HTTP server.




Figure 8. Configuration for two computers with a dedicated document file/HTTP server

Configuration of attached documents for multiple computers and multiple Hypertext Transfer Protocol servers

The multiple computers, multiple Hypertext Transfer Protocol (HTTP) servers scenario is applicable to both the WebSphere Application Server platform, and the WebLogic Server platform.

Table 58. Multiple computers, multiple HTTP servers, on Windows or UNIX

Operating system	Configuration	Specifications
Windows	- You store document files on a computer other than the application server computer that runs the system. - You store the document files for each Attached Documents folder on a different computer. - An HTTP server (such as Apache or Microsoft Internet Information Services) is on each computer that stores the document files. - For each folder in the system, you map a drive on the application server to point to the drive on the corresponding Document File/ HTTP server. The Document File/HTTP server is the computer that runs the HTTP server and stores the documents.	- Three HTTP server computers store document files: servers A, B, and C. - Server A stores the document files for the Attachments folder in the system, and document files for which no file path is specified. - Server B stores the document files for the Diagrams folder. - Server C stores the document files for the Images folder. - D is the drive on each HTTP server computer that stores the documents. - H, I, and J are mapped drives on the application server computer that runs the system. These drives correspond to Drive D on the HTTP server computers A, B, and C. - File names are case-sensitive.
UNIX	- You store document files on computers other than the application server computer that runs the system. - You store the document files for each Attached Documents folder in the system on a different computer. - An HTTP server (such as Apache or any other Web server) is on each computer that stores the document files. - You mount the network file system that contains the document files from the document file server onto the application server computer (UNIX only).	- Three HTTP server computers store document files: computers A, B, and C. - Server A stores the document files for the Attachments folder in the system, and the document files for which no file path is specified. - Server B stores the document files for the Diagrams folder. - Server C stores the document files for the Images folder. /d01, /d02, and /d03 are the NFS mount points on the application server computer for the home/file system on each of the HTTP servers. - Drive letter, file names, and directory names are case-sensitive.

The following graphic depicts a multiple computer configuration with multiple dedicated document file/HTTP servers.




Figure 9. Multiple computer configuration with multiple dedicated document file/HTTP servers

Alternative configurations for attached documents

There are several alternative configurations for the Attached Documents action when you are using either WebLogic Server or WebSphere Application Server platform.

- Two Computers, Local HTTP Server - Windows or UNIX
 - Store document files on a different computer than the application server computer.
 - The document HTTP server is on the application server computer that runs the system.

This configuration is for WebSphere Application Server only.

- Two Computers, One Dedicated HTTP Server - Windows or UNIX
 - Store document files on a different computer than the application server computer that runs the system.

- The HTTP server is on the computer that stores the document files.
- Multiple Computers, Multiple HTTP Servers - Windows or UNIX
 - Store document files on different servers, with each folder associated with a different server (and possibly managed by a different group). For example, store diagrams, images, and attachments on separate servers.
 - Each system that stores documents has its own HTTP server.

Multi-purpose internet mail extension mappings for WebLogic Server

Multi-Purpose Internet Mail Extensions (MIME) mapping associates a file name extension with a data file type (text, audio, image). You use these properties to map a MIME type to a file name extension. MIME is only for WebLogic Server.

The MIME-mapping element in a `web.xml` file defines the mapping between a file name extension and a MIME type. When you create a `doclinks\WEB-INF` directory, you copy a `web.xml` file into the directory. If you have trouble viewing certain document file types in the directory, review these steps:

- If you changed the `web.xml` file (or if you cannot open some attached documents before copying this file):
 1. Access Internet Explorer.
 2. Select **Tools/Internet Options**.
 3. On the **General** tab, under **Temporary Internet Files**, delete **Cookies** and delete **Files**.

Your browser might not display some document types (such as CAD diagrams) without special plug-ins. If you have these documents, check with your vendor to find out which plug-ins you need and if you can download them. If necessary, install the plug-ins on each client computer that is used to view and print these attached documents.

- If you have difficulty viewing certain types of documents, look at the mime-mapping sections of the `web.xml` file.

The `web.xml` file contains a series of parameters for mapping MIME data types. These parameters correspond to various types of document applications. For example, there is a parameter for `.doc` documents that corresponds to Microsoft Word documents:

- `<mime-mapping>`
- `<extensions>`
- `doc`
- `</extensions>`
- `<mime-type>`
- `application/msword`
- `</mime-type>`
- `</mime-mapping>`

The `<extension>` value is `doc`. and the `<mime-type>` value is `application/msword`.

The `web.xml` file can accommodate most common file types. If you have other types of documents that do not open for viewing as attachments, edit this file as follows to map your data type:

1. Copy a mime-mapping section in the file.
2. Paste it in a new section.

3. Change the appropriate application parameter lines to see the relevant applications extension and MIME type.

Finding the MIME type for an application

To find the MIME type for an application:

1. Access the Window registry.
2. Click **Windows Start** then click **Run** and type `regedit`.
3. Go to the `HKEY_CLASSES_ROOT` folder.
4. Expand the folder, and click the application extension. The mime-type appears on the Content Type line, under Data.

For example, for PDF documents, the MIME type is `application/pdf`.

5. After you edit the `web.xml` file, rebuild and redeploy the EAR file.

Working with attached documents

When you configure the **Attached Documents** action, you integrate the location of a stored document file with the location that you specified in the system. You can configure the system to store attached document files on the same server as the application server that is running the system. You can store attached document files on other servers as well.

Managing document libraries

You can store documents in an electronic document library that is located on a local server or remote server. Once you store a document in the library, it is possible to attach that document to records. You can also modify existing documents, and add URLs to the library.

Adding file attachments to the library

This action lets you attach a file attachment to a record (or to a task on a record).

Procedure

1. Open an application that has the **Attached Documents** action.
2. From the **Select Action** menu, navigate to **Attachment Library/Folders**, then select **Manage Library**.
3. Click **Add a Document to the Library**, then select **Add New File**.
4. Complete all necessary fields, and click **OK**.

Related concepts:

“Attached Documents overview” on page 307

You use the **Attached Documents** action, found in most system applications, to create a document library and to organize documents into folders. The system includes default folders. You can also create more folders or organize the folders into functional categories.

“Multi-purpose internet mail extension mappings for WebLogic Server” on page 313

Multi-Purpose Internet Mail Extensions (MIME) mapping associates a file name extension with a data file type (text, audio, image). You use these properties to map a MIME type to a file name extension. MIME is only for WebLogic Server.

Adding URLs to the library

This action lets you attach a URL to a record (or to a task on a record).

Procedure

1. Open an application that has the **Attached Documents** action.
2. From the **Select Action** menu, navigate to **Attachment Library/Folders**, then select **Manage Library**.
3. Click **Add a Document to the Library**, then select **Add New Web Page**.
4. Complete all necessary fields, and click **OK**.

Related concepts:

“Attached Documents overview” on page 307

You use the **Attached Documents** action, found in most system applications, to create a document library and to organize documents into folders. The system includes default folders. You can also create more folders or organize the folders into functional categories.

“Multi-purpose internet mail extension mappings for WebLogic Server” on page 313

Multi-Purpose Internet Mail Extensions (MIME) mapping associates a file name extension with a data file type (text, audio, image). You use these properties to map a MIME type to a file name extension. MIME is only for WebLogic Server.

Modifying existing documents

You can modify the information for documents that are stored in the document library.

About this task

If you click **View Details** for a document, you can modify only the **Document Description** field and the **URL/File Name** field, and select the **Print with Work Pack** check box.

Procedure

1. Open an application that has the **Attached Documents** action.
2. From the **Select Action** menu, navigate to **Attachment Library/Folders**, then select **Manage Library**.
3. In the Manage Library window, click the name of the document that you want to modify.
4. Modify editable fields that you want to change. Fields with a white background are editable and fields with a gray background are read-only.
5. Click **OK**.

Related concepts:

“Attached Documents overview” on page 307

You use the **Attached Documents** action, found in most system applications, to create a document library and to organize documents into folders. The system includes default folders. You can also create more folders or organize the folders into functional categories.

“Multi-purpose internet mail extension mappings for WebLogic Server” on page 313

Multi-Purpose Internet Mail Extensions (MIME) mapping associates a file name extension with a data file type (text, audio, image). You use these properties to map a MIME type to a file name extension. MIME is only for WebLogic Server.

Attaching documents to records

You can attach documents to records from within the library or from outside the library (and to task rows within a record).

Procedure

1. In an application that has the **Attached Documents** action, click **Attachments**.
2. Select one of the following actions to attach documents:

Option	Description
From within the library	Select Add from Library .
From outside the library	Select Add New File or select Add New Web Page .

3. If you added a file from outside the library, add the document to the library.
In either the Create a File Attachment window or the Create a URL Attachment window, you can select the **Add document to the document library for others to use** check box.

Related concepts:

“Attached Documents overview” on page 307

You use the **Attached Documents** action, found in most system applications, to create a document library and to organize documents into folders. The system includes default folders. You can also create more folders or organize the folders into functional categories.

“Multi-purpose internet mail extension mappings for WebLogic Server” on page 313

Multi-Purpose Internet Mail Extensions (MIME) mapping associates a file name extension with a data file type (text, audio, image). You use these properties to map a MIME type to a file name extension. MIME is only for WebLogic Server.

Printing work packs in a UNIX environment

You can print work packs in a UNIX environment.

Procedure

1. From the **Tools** menu in Microsoft Internet Explorer, select **Internet Options**.
2. On the **Security** tab, click **Custom Level**.
3. Under the **Initialize and script ActiveX controls not marked as safe setting**, click **Enable**.
4. Click **OK** to return to the **Security** tab, and click **OK** again.

Maintaining document libraries

Administrators maintain the library, create folders as needed, and specify the folders available for each application.

Procedure

1. Copy the file to the Attached Documents repository.
2. Specify a network path to the file, then attach the copy or the link to record.

Adding document folders

When you add a document folder, the folder is associated with the application to which you added it. Users can associate existing document folders with the current application, and add attachments to these folders.

Before you begin

You must have administrator privileges to access this action.

Procedure

1. Open any application that has the **Attached Documents** action.
2. From the **Select Action** menu, navigate to **Attachment Library/Folders** and select **Manage Folders**.
3. Click **Add a New Document Folder**.
4. Specify the information for the new folder.
5. Click **OK**.

Related concepts:

“Attached Documents overview” on page 307

You use the **Attached Documents** action, found in most system applications, to create a document library and to organize documents into folders. The system includes default folders. You can also create more folders or organize the folders into functional categories.

“Multi-purpose internet mail extension mappings for WebLogic Server” on page 313

Multi-Purpose Internet Mail Extensions (MIME) mapping associates a file name extension with a data file type (text, audio, image). You use these properties to map a MIME type to a file name extension. MIME is only for WebLogic Server.

Associating document folders with applications

You must associate document folders with an application before you can attach documents in those folders. By default, the Attachments folder, the Images folder, and the Diagrams folder are included in every application that has the **Attached Documents** action.

Before you begin

You must have administrator privileges to access this action.

Procedure

1. Open an application that has the **Attached Documents** action.
2. From the **Select Action** menu, navigate to **Attachment Library/Folders**, then select **Associate Folders**.
3. Click **New Row**.
4. In the **Document Folder** field, enter a value. The **Document Folder Description** field and the **Application** field contain default values, which you can change.
5. Click **OK** to save your changes.

Related concepts:

“Attached Documents overview” on page 307

You use the **Attached Documents** action, found in most system applications, to create a document library and to organize documents into folders. The system includes default folders. You can also create more folders or organize the folders into functional categories.

“Multi-purpose internet mail extension mappings for WebLogic Server” on page 313

Multi-Purpose Internet Mail Extensions (MIME) mapping associates a file name extension with a data file type (text, audio, image). You use these properties to map a MIME type to a file name extension. MIME is only for WebLogic Server.

Configuring attached documents in a single computer environment

Configuring WebLogic Server or WebSphere Application Server for the attached documents action requires two distinct tasks. If you are running WebLogic Server, you need to create a Web application. If you are running WebSphere Application Server, you need to edit the `httpd.conf`.

Creating attached documents directories in a single-computer environment

The first step in configuring attached documents is to create attached document directories for storing files. This task details the procedure for both WebSphere Application Server and WebLogic Server.

Procedure

1. Create a `doclinks` directory on the computer where the document files are stored. For example:

Operating system	Doclinks directory
Windows	<code>c:\doclinks</code>
UNIX	<code>/home/doclinks</code>

2. Share the drive so that users can connect to it.
3. Create the following subdirectories under the `doclinks` directory:
 - `attachments`
 - `default`
 - `diagrams`
 - `images`
4. Depending on which application server you are using, complete the steps for either WebSphere Application Server or a WebLogic Server:

Option	Description
If you used the WebSphere Application Server	Verify that the subdirectories were created as described in Step 3, and stop here.
If you used the WebLogic Server	Continue to Step 5.

5. Create another directory named `WEB-INF`.
6. Copy the `web.xml` file from the deployment folder into the directory that you created in Step 5:

Option	Description
Windows	<i>c:\install_home\deployment</i>
UNIX	<i>install_home/deployment</i>

The system contains other `web.xml` files. Be sure to copy the correct one. The file contains information for mapping MIME objects to customize.

- Verify that the subdirectories were created as described in Step 3 and Step 5.

Related concepts:

“Configuration of attached documents for a single computer” on page 308

For WebSphere Application Server and WebLogic Server, the single computer has a specific configuration and specifications. This configuration assumes that your computer is running a Windows or UNIX scenario.

Creating a Web application in a single-computer environment

If you are running WebLogic Server, creating a web application is the necessary process to configure the application server for attached documents.

Procedure

- Stop the WebLogic Server.
- Back up the `config.xml` file in the domain in which you want to configure the Web application:

Operating system	Path
Windows	<code><BEA WebLogic root>user_projects\domains\<domain name></code> For example, <code>/usr/BEA/user_projects/domains/mydomain</code> .
UNIX	<code><BEA WebLogic root> /user_projects/domains/<domain name>\</code> For example, <code>/usr/BEA/user_projects/domains/mydomain</code>

- Start the application server.
- Access the administration console:
`http://<hostname>:<port>/console`
 where `<hostname>` is the name of the computer and `<port>` is the port number of the application server.
- In the left pane under the Deployments node, click **Web Application Modules**.
- Delete the existing Web application named `doclinks`, if one exists on your system.
- In the right pane, click **Deploy a new Web Application Module**.
- Go to the `doclinks` directory and select it.
- Click **Target Module** at the bottom of the window.
- If you have more than one server, select the server on which you want to deploy your new Web application module.
 The name of the directory must be the root directory in which the documents are stored. Since you selected it in Step 7, `doclinks` is the default.
- Click **Continue**.

12. Click **Deploy**.

The Web application that you created appears in the Web Application tree in the left pane.

13. Verify that the doclinks Web module was installed correctly:

a. Complete one of the following steps:

- For Windows, create a test file, named test.txt, in the doclinks folder:
C:\doclinks\test.txt.
- For UNIX, create a test file, test.txt, at this location:
/home/doclinks/test.txt

b. Open a browser session and type the following address:

http://<server_name or ip address>:<port number>/doclinks/test.txt

For example:

http://localhost:7001/doclinks/test.txt

You can see your test.txt document in this window. If you cannot open the file, you did not correctly create the doclink web application. To reconfigure the doclink web application, restart these steps from the beginning.

What to do next

After you modify the location of the doclinks directory, edit the specified file paths in the system.

Related concepts:

“Configuration of attached documents for a single computer” on page 308

For WebSphere Application Server and WebLogic Server, the single computer has a specific configuration and specifications. This configuration assumes that your computer is running a Windows or UNIX scenario.

Editing the httpd.conf file in a single-computer environment

In WebSphere Application Server, the Attached Documents action uses the IBM HTTP Server to display attached documents. You must edit the httpd.conf file to specify the root of the \doclinks folder to be the home directory of WebSphere Application Server.

Procedure

1. Go to the location of the httpd.conf file for the IBM HTTP Server. The default installation location depends on your operating system:

Operating system	Path
Windows	C:\IBM HTTP Server\conf\httpd.conf
UNIX	/home/IBMHTTPD/conf/httpd.conf

2. Back up the httpd.conf file.
3. Open the httpd.conf file in a text editor. Find the section that begins with the following line:
This should be changed to whatever you set DocumentRoot to.
4. Change the Directory line you located in the previous step to specify the doclinks directory that you created:

Operating system	Directory line
Windows	<Directory C:\doclinks>
UNIX	<Directory /home/doclinks>

5. Find the section that begins with the following lines:

```
#
# Document Root: The directory out of which you will serve your
# documents. By default, all requests are taken from this directory, but
# symbolic links and aliases may be used to point to other locations.
#
```

6. Edit the DocumentRoot line to specify the doclinks directory that you created:

Operating system	Directory line
Windows	DocumentRoot C:\doclinks
UNIX	DocumentRoot /home/doclinks

7. Save and close the file.

8. Restart the HTTP server.

9. To verify that the HTTP server is configured correctly, complete the following steps:

a. Perform one of the following steps:

- For Windows, create a test file, test.txt, at this location:
C:\doclinks\test.txt
- For UNIX, create a test file, named test.txt, in the doclinks folder.
/home/doclinks/test.txt

b. Open a browser session and type the following address:

http://<server_name or ip address>/test.txt

For example:

http://localhost/test.txt

You can see your test.txt document in this window. If you cannot open the file, you must reconfigure the IBM HTTP Server. To reconfigure the IBM HTTP Server, repeat all the preceding steps.

10. Restart WebSphere Application Server and the system.

Related concepts:

“Configuration of attached documents for a single computer” on page 308

For WebSphere Application Server and WebLogic Server, the single computer has a specific configuration and specifications. This configuration assumes that your computer is running a Windows or UNIX scenario.

Editing default file paths in System Properties in a single-computer environment

After you modify the location of the doclinks directory, edit the specified file paths in the system. These steps are valid for both WebSphere Application Server and WebLogic Server.

Before you begin

You must have authorization to edit file paths in Attached Documents.

Procedure

1. Log in to the system and go to **System Configuration**. Select **Platform Configuration** and then **System Properties**.
2. Configure the Attached Documents properties as shown in the following table:

Property	Description	Global value
mxe.doclink.doctypes.defpath	The default file directory to use for folders in the library that do not have a default path specified in the database. The files for such folders are uploaded to the location.	• In Windows, WebLogic Server and WebSphere Application Server: C:\doclinks • In UNIX, WebLogic Server and WebSphere Application Server: /home/doclinks
mxe.doclink.maxfilesize	The maximum size (MB) for a file that you can upload to the Attached Documents Library folder.	Use the default value of 10 MB (10 = 10 MB) or replace it with a lesser value. Do not set the maximum file size to a value that exceeds the computer system capacity. If you do, a system error, OutOfMemory, occurs and the application server shuts down. To correct this error, change the value to less than 10 MB and restart the application server. If the value is set to 0, all file sizes are uploaded. However, there is the risk of OutOfMemory errors if a user uploads a large file size that exceeds system capacity. To correct this error, change the value to less than 10 MB and restart the application server.

Property	Description	Global value
mxe.doclink.path01	- The HTTP server path to link documents that are attached to records. Used to convert specified file paths of folders to URLs. - Use the following statement: <Value specified in the default path of a folder> = <URL from where the files are stored> The system reads the string, <Value specified in the default path of a folder>, and replaces it with the string, <URL from where the files are stored> For example, in Windows, the default file path for stored documents is C:\doclinks\diagrams. A user adds a document, diagram123.dwg, to the diagrams folder. The document is copied from the source to: C:\doclinks\diagrams. The mxe.doclink.path01 property converts the file path to http://localhost/doclinks/ diagrams. The link to view this file is http://localhost/doclinks/ diagrams/diagram123.dwg.	- Windows WebLogic Server: C<PATH>\doclinks=http://<servername or IP>:<port number>/doclinks For example C<PATH>\doclinks=http://localhost:7001/doclinks - UNIX WebLogic Server: /home/doclinks=http://<servername or IP>:<port number>/doclinks For example: /home/doclinks=http://localhost:7001/doclinks - Windows WebSphere Application Server: C<PATH>\doclinks=http://<servername or IP> For example: C<PATH>\doclinks=http://localhost - UNIX WebSphere Application Server: /home/doclinks = http://<servername or IP> For example: /home/doclinks = http://localhost/
mxe.doclink.multilang. aix.websphere	Indicates whether the application runs on AIX WebSphere Application Server platform. The default value is false.	- Change the value to true if it is running on AIX WebSphere Application Server platform. - Set the value to false if the application is running on other platforms, such as a system other than a WebSphere Application Server on AIX.

In the mxe.doclink.path01 property, the server name in the path must be a fully qualified server name.

3. Restart the application server.

Related concepts:

“Configuration of attached documents for a single computer” on page 308
For WebSphere Application Server and WebLogic Server, the single computer has a specific configuration and specifications. This configuration assumes that your computer is running a Windows or UNIX scenario.

Editing default file paths in System Properties for multiple computers and multiple Hypertext Transfer Protocol servers

Once you modify the location of the doclinks directory, you can use the Systems Properties application to edit the specified file paths in the system. This task is for both WebSphere Application Server and WebLogic Server, and for systems running Windows or UNIX operating systems.

Before you begin

You must have authorization to edit file paths in the **Attached Documents** action.

Procedure

1. Log in to the system and go to **System Configuration**. Select **Platform Configuration** and then **System Properties**.
2. Configure the following properties for the **Attached Documents** action:

Property	Description	Global value
mxe.doclink.doctypes.defpath	The default file directory to use for folders in the library that do not have a default path specified in the database. Files for these folders are uploaded to the location, mxe.doclink.doctypes.defpath.	• Windows: H:\doclinks • UNIX: /d01/doclinks
mxe.doclink.maxfilesize	The maximum size (in MB) for a file that you can upload to the Attached Documents Library folder.	Use the default value of 10 MB (10 = 10 MB) or replace it with a lesser value. Do not set the maximum file size to a value that exceeds the computer system capacity. If you do, a system error, OutOfMemory, occurs and the application server shuts down. To correct this error, change the value to less than 10 MB and restart the application server. If the value is set to 0, the system allows all file sizes to be uploaded. However, there is the risk of OutOfMemory errors if a user uploads a large file size that exceeds system capacity. To correct this error, change the value to less than 10 MB and restart the application server.

Property	Description	Global value
mxe.doclink.path01	- The HTTP server path to link documents that are attached records. - Use the following statement: <Value specified in the default path of a folder> = <URL from where the files are served> The system reads the string, <Value specified in the default path of a folder>, and replaces it with the string, <URL from where the files are served>. For example, in Windows, the default file path for stored documents is H:\doclinks\diagrams. A user adds a document, diagram123.dwg, to the diagrams folder. The document is copied from the source to: H:\doclinks\diagrams. The mxe.doclink.path01 property converts the file path to http://localhost/doclinks/diagrams. The link to view this file is http://localhost/doclinks/diagrams/diagram123.dwg.	- Windows: H<PATH>\doclinks=http://dochostA - UNIX: /d01/doclinks=http://dochostA
mxe.doclink.path02	The HTTP server path to link documents attached to system records. Used to convert specified file paths of folders to URLs.	- Windows: I<PATH>\doclinks=http://dochostB - UNIX: /d02/doclinks=http://dochostB
mxe.doclink.path03	The HTTP server path to link documents attached to records. Used to convert specified file paths of folders to URLs.	- Windows: J<PATH>\doclinks=http://dochostC - UNIX: /d03/doclinks=http://dochostC
mxe.doclink.multilang.aix.websphere	Indicate whether the application is running on AIX WebSphere Application Server platform. The default value is false.	- Change the value to true if it is running on AIX WebSphere Application Server platform. - Set the value to false if the application is running on other platforms, such as a system other than WebSphere Application Server on AIX.

Because multiple entries are allowed (by default, up to 10) to convert file paths, you can set up your system so that each document folder uses different servers or directories. In the mxe.doclink.path n property, the dochost in the path must be a fully qualified server name.

3. Restart the application server.

Related concepts:

“Configuration of attached documents for a single computer” on page 308
For WebSphere Application Server and WebLogic Server, the single computer has a specific configuration and specifications. This configuration assumes that your computer is running a Windows or UNIX scenario.

Changing paths for demo data library files in a single-computer environment

A demo attachment library, named DATA, is included with the Attached Documents action. To view these library files, running on either WebSphere Application Server or WebLogic Server, change the file to be the same as your doclinks directory setup.

Before you begin

You must have authorization to edit file paths in Attached Documents.

Procedure

1. Log in to the system and open an application that has the **Attached Documents** action.
2. From the **Select Action** menu, go to **Attachment Library/Folders** and select **Manage Library**.
3. In the Manage Library window, click the **Details** icon for the document whose file path you want to change.
4. In the **URL/File Name** field, specify the new location of the doclinks directory. Type the full path, including the drive letter.
The drive letter, path, and folder names are case-sensitive. They must be under the same as the path and folder names that you created in the System Properties application.
5. Change the file paths for each document to the following paths:

Operating system	File paths
Windows	C:\doclinks\<filename> For example, document 1001, URL/File Name is displayed as default as: \DOCLINKS\BOILDER.DWF Change it to: C:\doclinks\BOILDER.DWF
UNIX	home/doclinks/<filename> For example, document 1001, URL/File Name is displayed as default as: \DOCLINKS\BOILDER.DWF Change it to: /home/doclinks/BOILDER.DWF

You must change every library file path in the dialog that appears in the window.

6. Click **OK**.
7. Restart the application server.

Related concepts:

“Configuration of attached documents for a single computer” on page 308

For WebSphere Application Server and WebLogic Server, the single computer has a specific configuration and specifications. This configuration assumes that your computer is running a Windows or UNIX scenario.

Configuring attached documents for two computers and a local Hypertext Transfer Protocol server

When you use two computers, a local HTTP server, and WebSphere Application Server on either Windows or UNIX, certain configuration specifications apply.

Creating attached documents directories for two computers and a local Hypertext Transfer Protocol server

The first step in configuring your system for attached documents is to create directories to store files.

Procedure

1. Stop the WebSphere Application Server.
2. Create the following subdirectories under doclinks:
 - attachments
 - default
 - diagrams
 - images
3. Verify the directory structure, and complete the following steps:
 - a. Create another directory named WEB-INF.
 - b. Go to the doclinks directory created in Step 1.
 - c. Copy the web.xml file from the deployment folder into the WEB-INF directory that you created.

The system contains several additional web.xml files. Be sure to copy the correct one.

- For Windows, the web.xml file is: `c:\install_home\deployment`
- For UNIX, the web.xml file is: `install_home/deployment`

The file contains mime-mapping information that you can customize.

- d. Verify the directory structure.

Example

The directories are found in the following locations:

Operating system	Doclinks directory
Windows	D:\doclinks
UNIX	/home/doclinks

Related concepts:

“Configuration of attached documents for two computers and a local Hypertext Transfer Protocol server” on page 308

When you use two computers, a local Hypertext Transfer Protocol (HTTP) server, and WebSphere Application Server on either Windows or UNIX, certain configuration specifications apply.

Creating Web applications for two computers and a local Hypertext Transfer Protocol server

If you are running WebLogic Server, creating a web application is the necessary process to configure the application server for attached documents.

Procedure

1. Stop the WebLogic Server.
2. Back up the config.xml file in the domain in which you want to configure the Web application.

Operating system	Path
Windows	<BEA WebLogic root>\user_projects\domains\<domain_name> For example, \usr\bea\user_projects\domains\mydomain
UNIX	<BEA WebLogic root>/user_projects/domains/<domain_name> For example, /usr/bea/user_projects/domains/mydomain

3. Start the WebLogic Server.
4. To log in to the administration console, type the following URL:
http://<hostname>:<port>/console
where <hostname> is the name of the computer and <port> is the port number of the application server.
5. In the left pane, under the Deployments node, click **Web Application Modules**.
6. Delete the existing Web application named doclinks, if one exists on your system.
7. In the right pane, click **Deploy a new Web Application Module**.
8. Go to the location of the doclinks directory on the mapped drive.

Operating system	Doclinks directory location
Windows	For example: 1. Click the computer name to display the drive letters. 2. Click the mapped drive, H, to display the directories on H (which is the drive D on the computer that stores the document files). The doclinks directory that you created on D appears in the list following the path statement.

Operating system	Doclinks directory location
UNIX	For example: 1. Click the host name to display the root file system. 2. Click /d01 to display the directories that /d01 references on the computer that stores the document files. The doclinks directory that you created is displayed in the list following the path statement.

9. Select the doclinks directory.
10. Click **Target Module** at the bottom of the screen.
11. If you have more than one server, select the server on which you want to deploy your new Web application module, then click **Continue**.
12. Review your choices.

The name must be the root directory name where the documents are stored. Since you selected the doclinks directory in Step 9, doclinks is the default. The name is case-sensitive.
13. Click **Deploy**.

The Web application that you created appears in the Web Application tree in the left pane.
14. Complete the following steps to verify that the doclinks web module was installed correctly:
 - a. Depending on whether you use Windows or UNIX, complete one of the following steps:
 - For Windows, create a test file, named test.txt, at this location:
D:\doclinks\test.txt
 - For UNIX, create a test file, named test.txt, at this location:
/home/doclinks/test.txt
 - b. Open a browser session and type the following address:
http://<server_name or ip address>:<port number>/doclinks/test.txt
For example:
http://localhost:7001/doclinks/test.txt

You can see your test.txt document. If you cannot open the file, you did not correctly configure WebLogic Server for doclinks. To create the doclink web application. To reconfigure the doclink web application, redo this task.

Related concepts:

“Configuration of attached documents for two computers and a local Hypertext Transfer Protocol server” on page 308

When you use two computers, a local Hypertext Transfer Protocol (HTTP) server, and WebSphere Application Server on either Windows or UNIX, certain configuration specifications apply.

Editing default file paths in System Properties for two computers and a local Hypertext Transfer Protocol server

Because you modified the location of the doclinks directory, you can use the Systems Properties application to edit the specified file paths. These steps are valid for systems running on Windows or UNIX.

Before you begin

You must have authorization to edit file paths in the **Attached Documents** action.

Procedure

1. Log in to the system and go to **System Configuration**. Select **Platform Configuration** and then click **System Properties**.
2. Configure the **Attached Documents** action as shown in the following table.

Property	Description	Global value
mxe.doclink.doctypes.defpath	The default file directory to use for folders in the library that do not have a default path specified in the database.	Windows: H:\doclinks UNIX: /d01/doclinks
mxe.doclink.maxfilesize	The maximum size (in MB) for a file that you can upload to the Attached Documents Library folder.	Use the default value of 10 MB (10 = 10 MB) or replace it with a lesser value. Do not set the maximum file size to a value that exceeds the computer system capacity. If you do, a system error, OutOfMemory, occurs and the application server shuts down. To correct this error, change the value to less than 10 MB and restart the application server. If the value is set to 0, the system allows all file sizes to be uploaded. However, there is the risk of OutOfMemory errors if a user uploads a large file size that exceeds system capacity. To correct this error, change the value to less than 10 MB and restart the application server.

Property	Description	Global value
mxe.doclink.path01	The HTTP server path to link documents that are attached to records. Used to convert specified file paths of folders to URLs. Use the following statement: <Value specified in the default path of a folder> = <URL from where the files are served> The system reads the string, <Value specified in the default path of a folder>, and replaces it with the string, <URL from where the files are served> For example, in Windows, the default file path for stored documents is H:\doclinks\diagrams. A user adds a document, diagram123.dwg, to the diagrams folder. The document is copied from the source to: H:\doclinks\diagrams. The mxe.doclink.path01 property converts the file path to http://localhost/doclinks/diagrams. The link to view this file is http://localhost/doclinks/ diagrams/ diagram123.dwg.	Windows: H<PATH>\doclinks=http://<servername or IP>:<port number>/doclinks For example: H<PATH>\doclinks=http://localhost:7001/doclinks UNIX: /d01/doclinks=http://<servername or IP>:<port number>/doclinks For example: /d01/doclinks=http://localhost:7001/doclinks

In the mxe.doclink.pathn01 property, the server name in the path must be a fully qualified server name.

3. Restart the application server.

Related concepts:

“Configuration of attached documents for two computers and a local Hypertext Transfer Protocol server” on page 308

When you use two computers, a local Hypertext Transfer Protocol (HTTP) server, and WebSphere Application Server on either Windows or UNIX, certain configuration specifications apply.

Editing default file paths in related applications for two computers and a local Hypertext Transfer Protocol server

Because you modified the location of the doclinks directory, you then edit the specified file paths in the system. Complete the following steps in an application that has the **Attached Documents** action. These steps are valid for systems running on Windows or UNIX.

Before you begin

You be authorized to edit file paths in the **Attached Documents** action.

Procedure

1. Log in to the system and open an application that has the Attached Documents action.
2. From the **Select Action** menu, go to **Attachment Library/Folders** and select **Manage Folders**.
3. In the Manage All Documents Folder window, click the **Details** icon next to the document folder whose file path you want to change.
4. In the **Default File Path** field, edit the path to specify the new location of the associated directory. Type the full path using the mapped drive letter.

The drive letter, path, and folder name are case sensitive and must be under the same path and folder names that you created earlier.

Change the file paths for the Attachments, CAD, Diagrams, and Images folders to the following file paths:

Operating system	File paths
Windows	H:\doclinks\attachments H:\doclinks\cad H:\doclinks\diagrams H:\doclinks\images
UNIX	/d01/doclinks/attachments /d01/doclinks/cad /d01/doclinks/diagrams /d01/doclinks/images

If you create additional attached document folders, you also edit their file paths.

5. Click **OK**.
6. Restart the application server.

Related concepts:

“Configuration of attached documents for two computers and a local Hypertext Transfer Protocol server” on page 308

When you use two computers, a local Hypertext Transfer Protocol (HTTP) server, and WebSphere Application Server on either Windows or UNIX, certain configuration specifications apply.

Changing paths for demo data library files for two computers and a local Hypertext Transfer Protocol server

A demo attachment library, named DATA, is included with the **Attached Documents** action. To view these library files when you are using a WebLogic server platform, change the file paths to be the same as your doclinks directory setup. These steps are valid for systems running on Windows or UNIX.

Before you begin

You must have authorization to edit file paths in the **Attached Documents** action.

Procedure

1. Log in to the system and open an application that has the Attached Documents action.

- From the **Select Action** menu, go to **Attachment Library/Folders** and select **Manage Library**.
- In the Manage Library window, click the **Details** icon next to the document whose file path you want to change.
- In the **URL/File Name** field, change the path to specify the new location of the doclinks directory. Type the full path and use the mapped drive letter.
The drive letter, path, and folder names are case-sensitive, and must under the same the path and folder names that you created earlier.
- Change the file paths for each document:

Operating system	File paths
Windows	H:\doclinks\<filename> For example, document 1001, URL/File Name is displayed as default as: \DOCLINKS\BOILER.DWF Change it to: H:\doclinks\BOILER.DFW
UNIX	/d01/doclinks/<filename> For example, document 1001, URL/File Name is displayed as default as: \DOCLINKS\BOILER.DWF Change it to: /d01/doclinks/BOILER.DFW

You must modify every listed library file path in the window.

- Click **OK**.
- Restart the WebLogic Server.

Related concepts:

“Configuration of attached documents for two computers and a local Hypertext Transfer Protocol server” on page 308

When you use two computers, a local Hypertext Transfer Protocol (HTTP) server, and WebSphere Application Server on either Windows or UNIX, certain configuration specifications apply.

Configuring attached documents for two computers and a dedicated Hypertext Transfer Protocol server

When you use two computers with a dedicated HTTP server on either Windows or UNIX, specific configuration steps and specifications apply.

Creating attached documents directories for two computers and a dedicated Hypertext Transfer Protocol server

The first step in configuring for attached documents is to create directories to store the files. This task is for both WebSphere Application Server and WebLogic Server and is applicable for systems running either Windows or UNIX.

Procedure

- Create a doclinks directory on the computer that stores the document files.
For example:

Operating system	Doclinks directory
Windows	D:\doclinks

Operating system	Doclinks directory
UNIX	/home/doclinks

2. Create the following subdirectories under doclinks:

- attachments
- default
- diagrams
- images

If you created additional attached document folders, then create subdirectories for them.

3. On the application server computer that runs the system, perform one of these tasks:

Option	Description
Windows	Map drive H to drive D on the computer on which the documents are stored.
UNIX	Configure /d01 to be the NFS mount point for the / home file system on the HTTP server that stores the document files.

Related concepts:

“Configuration of attached documents for two computers and a dedicated Hypertext Transfer Protocol server” on page 309

When you use two computers with a dedicated Hypertext Transfer Protocol (HTTP) server on either Windows or UNIX, the following configuration and specifications apply. This configuration applies to the WebSphere Application Server platform or to the WebLogic Server platform.

Setting up the server for attached documents for two computers and a dedicated Hypertext Transfer Protocol server

This configuration scenario relies on an HTTP server that is independent of the system. You can use the HTTP server application that you prefer, either WebSphere Application Server or WebLogic Server.

Procedure

1. Depending on your operating system, perform one of the following tasks:

Option	Description
Windows	In Apache, edit the httpd.conf file to use d:\doclinks as its default home page documents directory.
UNIX	In Apache, edit the httpd.conf file to use /home/doclinks as its default home page documents directory.

2. Because you edited the httpd.conf file, restart the HTTP server.

Related concepts:

“Configuration of attached documents for two computers and a dedicated Hypertext Transfer Protocol server” on page 309

When you use two computers with a dedicated Hypertext Transfer Protocol (HTTP) server on either Windows or UNIX, the following configuration and specifications apply. This configuration applies to the WebSphere Application Server platform or to the WebLogic Server platform.

Editing default file paths in System Properties for two computers and a dedicated HTTP server

After you modify the location of the doclinks directory, you can use the Systems Properties application to edit the specified file paths in the system. These steps are for WebSphere Application Server and WebLogic Server, and systems running on Windows or UNIX.

Before you begin

You must have authorization to edit file paths in the **Attached Documents** action.

Procedure

1. Log in to the system and go to **System Configuration**. Select **Platform Configuration** then **System Properties**.
2. Configure the Attached Documents actions is described in the following table:

Property	Description	Global value
mxe.doclink.doctypes.defpath	The default file directory to use for folders in the library that do not have a default path specified in the database.	Windows:H:\doclinks UNIX: /d01/doclinks
mxe.doclink.maxfilesize	The maximum size (in MB) for a file that you can upload to the Attached Documents Library folder.	Use the default value of 10 MB (10 = 10 MB) or replace it with a lesser value. Do not set the maximum file size to a value that exceeds the computer system capacity. If you do, a system error, <i>OutOfMemory</i> , occurs and the application server shuts down. To correct this error, change the value to less than 10 MB and restart the application server. If the value is set to 0, the system allows all file sizes to be uploaded. However, there is the risk of <i>OutOfMemory</i> errors if a user uploads a large file size that exceeds system capacity. To correct this error, change the value to less than 10 MB and restart the application server.

Property	Description	Global value
mxe.doclink.path01	The HTTP server path to link documents that are attached records. Used to convert specified file paths of folders to URLs. Use the following statement: <Value specified in the default path of a folder> = <URL from where the files are served> The system reads the string, <Value specified in the default path of a folder>, and replaces it with the string, <URL from where the files are served>. For example, in Windows, the default file path for stored documents is H:\doclinks\diagrams. A user adds a document, diagram123.dwg, to the diagrams folder. The document is copied from the source to: H:\doclinks\diagrams. The mxe.doclink.path01 property converts the file path to http://localhost/doclinks/diagrams. The link to view this file is http://localhost/doclinks/diagrams/ diagram123.dwg.	Windows: H<PATH>\doclinks=http://dochost/ UNIX: /d01/ doclinks=http://dochost
mxe.doclink.multilang. aix.websphere	Indicate whether the application is running on AIX WebSphere Application Server platform. Default value is false.	Change the value to true if it is running on AIX WebSphere Application Server platform. Set the value to false if the application is running on other platforms, such as a system other than WebSphere Application Server on AIX.

In the mxe.doclink.path01 property, the dochost in the path must be a fully qualified server name.

- Restart the application server.

Related concepts:

“Configuration of attached documents for two computers and a dedicated Hypertext Transfer Protocol server” on page 309

When you use two computers with a dedicated Hypertext Transfer Protocol (HTTP) server on either Windows or UNIX, the following configuration and specifications apply. This configuration applies to the WebSphere Application Server platform or to the WebLogic Server platform.

Editing default file paths in related applications for two computers and a dedicated Hypertext Transfer Protocol server

Once you modify the location of the doclinks directory, you can edit the specified file paths in the system. To edit default file paths, complete the following steps in any application that uses the Attached Documents action. These steps are for systems running WebSphere Application Server or WebLogic Server.

Before you begin

You must have authorization to edit file paths in the **Attached Documents** action.

Procedure

1. Log in to the system and open an application that has the **Attached Documents** action.
2. From the **Select Action** menu, go to **Attachment Library/Folders** and then select **Manage Folders**.
3. In the Manage All Document folders window, click the **Details** icon next to the document folder whose file path you want to change.
4. In the **Default File Path** field, edit the path to specify the new location of the associated directory. Type the full path using the mapped drive letter.

The drive letter, path, and folder names are case-sensitive and must be under the same path and folder names that you created earlier.

Change the file paths for the Attachments, CAD, Diagrams, and Images folders to:

Operating system	File paths
Windows	H:\doclinks\attachments H:\doclinks\cad H:\doclinks\diagrams H:\doclinks\images
UNIX	/d01/doclinks/attachments /d01/doclinks/cad /d01/doclinks/diagrams /d01/doclinks/images

If you create additional attached document folders, you also edit their file paths.

5. Click **OK**.
6. Restart the application server.

Related concepts:

“Configuration of attached documents for two computers and a dedicated Hypertext Transfer Protocol server” on page 309

When you use two computers with a dedicated Hypertext Transfer Protocol (HTTP) server on either Windows or UNIX, the following configuration and specifications apply. This configuration applies to the WebSphere Application Server platform or to the WebLogic Server platform.

Changing paths for demo data library files for two computers and a dedicated Hypertext Transfer Protocol server

A demo attachment library, named DATA, is included with the Attached Documents action. To view these library files from the system running on the WebSphere Application Server platform or the WebLogic Server platform, change the file paths to be the same as your doclinks directory setup.

Before you begin

You must have authorization to edit file paths in the **Attached Documents** action.

Procedure

1. Log in to the system and open an application that has the **Attached Documents** action.
2. From the **Select Action** menu, go to **Attachment Library/Folders** and then select **Manage Library**.
3. In the Manage Library window, click the **Details** icon next to the document whose file path you want to change.
4. In the **URL/File Name** field, edit the path to specify the new location of the doclinks directory. Type the full path using the mapped drive letter.

The drive letter, path, and folder names are case-sensitive and must be under the same path and folder names that you created earlier.

Change the file paths for each document to the following file paths:

Operating system	File paths
Windows	H:\doclinks\<filename> For example, document 1001, URL/File Name is displayed as default as: \\DOCLINKS\BOILDER.DWF Change it to: H:\doclinks\BOILDER.DFW
UNIX	/d01/doclinks/<filename> For example, document 1001, URL/File Name is displayed as default as: \\DOCLINKS\BOILDER.DWF Change it to: /d01/doclinks/BOILDER.DFW

You must modify every listed library file path in the dialog box.

5. Click **OK**.
6. Perform one of the following steps:

Option	Description
If you are using WebSphere Application Server and you edited the httpd.conf file	Restart the HTTP server, the WebSphere Application Server, and the system.
If you are using a WebLogic Server	Restart the application server.

Related concepts:

“Configuration of attached documents for two computers and a dedicated Hypertext Transfer Protocol server” on page 309

When you use two computers with a dedicated Hypertext Transfer Protocol (HTTP) server on either Windows or UNIX, the following configuration and specifications apply. This configuration applies to the WebSphere Application Server platform or to the WebLogic Server platform.

Configuring attached documents for multiple computers and multiple Hypertext Transfer Protocol servers

When you use multiple computers with multiple Hypertext Transfer Protocol (HTTP) servers on either Windows or UNIX, specific configuration steps and specifications apply.

Creating attached documents directories for multiple computers and multiple Hypertext Transfer Protocol servers

The first step in configuring for attached documents is to create directories to store the files. These steps are valid for both WebSphere and WebLogic.

Procedure

1. Create a doclinks directory on the computer that stores the document files.

For example:

Operating system	Doclinks directory
Windows	D:\doclinks
UNIX	/home/doclinks

2. Create the following subdirectories under doclinks for each server:

Operating system	Doclinks directory
Windows	Server A: doclinks\attachments
	Server A: doclinks\default
	Server B: doclinks\diagrams
	Server C: doclinks\images
UNIX	Server A: /home/doclinks/attachments
	Server A: /home/doclinks/default
	Server B: /home/doclinks/diagrams
	Server C: /home/doclinks/images

3. On the application server computer that runs the system, perform the following tasks to map the drives:

Operating system	Map drive
Windows	• Map drive H to drive D on server A. • Map drive I to drive D on server B. • Map drive J to drive D on server C.
UNIX	1. Configure /d01 to be the NFS mount point for the /home file system on server A. 2. Configure /d02 to be the NFS mount point for the /home file system on server B. 3. Configure /d03 to be the NFS mount point for the /home file system on server C.

Related concepts:

“Configuration of attached documents for multiple computers and multiple Hypertext Transfer Protocol servers” on page 310

The multiple computers, multiple Hypertext Transfer Protocol (HTTP) servers scenario is applicable to both the WebSphere Application Server platform, and the WebLogic Server platform.

Setting up the server for attached documents for multiple computers and multiple Hypertext Transfer Protocol servers

The multiple computer dedicated HTTP server scenario relies on HTTP servers that are independent of the system. You can use the HTTP server application that you prefer.

Procedure

1. Perform only one of the following tasks, depending on your operating system:

Option	Description
Windows	In Apache, edit the httpd.conf file to use d:\doclinks as its default home page documents directory.
Windows	In Microsoft Internet Information Services, create a virtual folder named doclinks and point it to the d:\doclinks directory on the same computer. You can also point the Microsoft Internet Information Services default home page directory to d:\doclinks.
UNIX	In Apache, edit the httpd.conf file to use /home/doclinks as its default home page documents directory.

2. Once you edit the httpd.conf file, restart the HTTP server.

Related concepts:

“Configuration of attached documents for multiple computers and multiple Hypertext Transfer Protocol servers” on page 310

The multiple computers, multiple Hypertext Transfer Protocol (HTTP) servers scenario is applicable to both the WebSphere Application Server platform, and the WebLogic Server platform.

Editing default file paths in System Properties for multiple computers and multiple Hypertext Transfer Protocol servers

Once you modify the location of the doclinks directory, you can use the Systems Properties application to edit the specified file paths in the system. This task is for both WebSphere Application Server and WebLogic Server, and for systems running Windows or UNIX.

Before you begin

You must have authorization to edit file paths in the **Attached Documents** action.

Procedure

1. Log in to the system and go to **System Configuration**. Select **Platform Configuration** and then **System Properties**.
2. Configure the following properties for the **Attached Documents** action:

Header	Header	Header
mxe.doclink.doctypes.defpath	The default file directory to use for folders in the library that do not have a default path specified in the database.	Windows: H:\doclinks UNIX: /d01/doclinks
mxe.doclink.maxfilesize	The maximum size (in MB) for a file that you can upload to the Attached Documents Library folder.	Use the default value of 10 MB (10 = 10 MB) or replace it with a lesser value. Do not set the maximum file size to a value that exceeds the computer system capacity. If you do, a system error, OutOfMemory, occurs and the application server shuts down. To correct this error, change the value to less than 10 MB and restart the application server. If the value is set to 0, the system allows all file sizes to be uploaded. However, there is the risk of OutOfMemory errors if a user uploads a large file size that exceeds system capacity. To correct this error, change the value to less than 10 MB and restart the application server.

Header	Header	Header
mx.e.doclink.path01	The HTTP server path to link documents that are attached records. Use the following statement: <Value specified in the default path of a folder> = <URL from where the files are served> The system reads the string, <Value specified in the default path of a folder>, and replaces it with the string, <URL from where the files are served>. For example, in Windows, the default file path for stored documents is H:\doclinks\diagrams. A user adds a document, diagram123.dwg, to the diagrams folder. The document is copied from the source to: H:\doclinks\diagrams. The mx.e.doclink.path01 property converts the file path to http://localhost/doclinks/diagrams. The link to view this file is http://localhost/doclinks/diagrams/diagram123.dwg.	Windows: H<PATH>\doclinks=http://dochostA UNIX: /d01/doclinks=http://dochostA
mx.e.doclink.path02	The HTTP server path to link documents attached to system records. Used to convert specified file paths of folders to URLs.	Windows: I<PATH>\doclinks=http://dochostB UNIX: /d02/doclinks=http://dochostB
mx.e.doclink.path03	The HTTP server path to link documents attached to records. Used to convert specified file paths of folders to URLs.	Windows: J<PATH>\doclinks=http://dochostC UNIX: /d03/doclinks=http://dochostC
mx.e.doclink.multilang.aix.web sphere	Indicate whether the application is running on AIX WebSphere Application Server platform. The default value is false.	Change the value to true if it is running on AIX WebSphere Application Server platform. Set the value to false if the application is running on other platforms, such as a system other than WebSphere Application Server on AIX.

Because multiple entries are allowed (by default, up to 10) to convert file paths, you can set up your system so that each document folder uses different servers or directories. In the mx.e.doclink.pathnn property, the dochost in the path must be a fully qualified server name.

3. Restart the application server.

Related concepts:

“Configuration of attached documents for multiple computers and multiple Hypertext Transfer Protocol servers” on page 310

The multiple computers, multiple Hypertext Transfer Protocol (HTTP) servers scenario is applicable to both the WebSphere Application Server platform, and the WebLogic Server platform.

Editing default file paths in related applications for multiple computers and multiple Hypertext Transfer Protocol servers

Once you modify the location of the doclinks directory, you can edit the specified file paths in the system. This task is for both WebSphere Application Server and WebLogic Server.

Before you begin

You must have authorization to edit file paths in the **Attached Documents** action.

Procedure

1. Log in to the system and open an application that has the **Attached Documents** action.
2. From the **Select Action** menu, select **Attachment Library/Folders** and then **Manage Folders**.
3. In the Manage All Documents Folders window, click the **Details** icon next to the document folder whose file path you want to change.
4. In the **Default File Path** field, edit the path to specify the new location of the associated directory. Type the full path using the mapped drive letter.

The drive letter, path, and folder names are case-sensitive and must under be the same path and folder names that you created earlier.

Change the file paths for the Attachments, CAD, Diagrams, and Images folders to the following file paths:

Operating system	File paths
Windows	H:\doclinks\attachments
	I:\doclinks\diagrams
	J:\doclinks\images
UNIX	/d01/doclinks/attachments
	/d02/doclinks/diagrams
	/d03/doclinks/images

If you create additional attached document folders, you can edit their file paths.

5. Click **OK**.
6. Restart the application server.

Related concepts:

“Configuration of attached documents for multiple computers and multiple Hypertext Transfer Protocol servers” on page 310

The multiple computers, multiple Hypertext Transfer Protocol (HTTP) servers scenario is applicable to both the WebSphere Application Server platform, and the WebLogic Server platform.

Changing paths for demo data library files for multiple computers and multiple Hypertext Transfer Protocol servers

A demo attachment library, named DATA, is included with the **Attached Documents** action. To view these library files from the system running on WebSphere Application Server platform or WebLogic Server platform, change the file path to be the same as your doclinks directory setup.

Before you begin

You must have authorization to edit file paths in the **Attached Documents** action.

Procedure

1. Log in to the system and open an application that has the **Attached Documents** action.
2. From the **Select Action** menu, go to **Attachment Library/Folders** and select **Manage Library**.
3. In the Manage Library window, click the **Details** icon next to the document whose file path you want to change.
4. In the URL/File Name field, edit the path to specify the new location of the doclinks directory. Type the full path using the mapped drive letter.
The drive letter, path, and folder names are case-sensitive and must be under the same path and folder names that you created earlier.
5. Change the file paths for each document to the following file paths:

Operating system	File paths
Windows	C:\doclinks\<filename> For example, document 1001, URL/File Name is displayed as default as: \DOCLINKS\BOILDER.DWF Change it to: C:\doclinks\BOILDER.DFW
UNIX	/d01/doclinks/<filename> For example, document 1001, URL/File Name is displayed as default as: \DOCLINKS\BOILDER.DWF Change it to: /home/doclinks/BOILDER.DFW

You must modify every listed library file path in the dialog box.

6. Click **OK**.
7. Perform one of the following steps:

Option	Description
If you are using WebSphere Application Server and you edited the httpd.conf file	Restart the HTTP server, the WebSphere Application Server, and the system.
If you are using a WebLogic Server	Restart the WebLogic Server.

Related concepts:

“Configuration of attached documents for multiple computers and multiple Hypertext Transfer Protocol servers” on page 310

The multiple computers, multiple Hypertext Transfer Protocol (HTTP) servers scenario is applicable to both the WebSphere Application Server platform, and the WebLogic Server platform.

Chapter 14. Managing log files

You create log files to save the informational, warning, or error messages about the system. You can also manage the logging process and the format of the log files across your organization.

Logging overview

You can create and manage log files that contain informational, warning, or error messages about the system.

Logging application components

The Logging application contains several components.

Logging has the following components:

- Loggers
- Appenders
- Layouts

Loggers

Loggers are components of the logging process that prepare log statements to be written to console or log file.

Loggers are named entities or keys, such as `log4j.logger.maximo.sql`. Loggers also form a hierarchy. A logger is defined as an ancestor of another logger. This relationship is true only when the name of the logger is followed by a dot or is a prefix of the descendant logger name. If there are no ancestors between a logger and the descendant logger, a logger becomes the parent of a child logger. For example, `log4j.logger.maximo.sql` is the parent of `log4j.logger.maximo.sql.WORKORDER`.

You can assign the following levels to Loggers: `DEBUG`, `INFO`, `WARN`, `ERROR`, and `FATAL`. A level indicates a type of event that the system logs.

Appenders

Appenders are components of the logging process. You can send logging requests to multiple destinations. These output destinations are called appenders.

Appenders can exist for consoles or files. You can associate one or more loggers with a given appender. Alternatively, you can associate a single logger with multiple appenders.

Appender types

The system provides you with the following types of appenders. You cannot delete any of the system appenders.

Table 59. Appender types

Type	Description
Console appender	Writes log statements to the application server console.

Table 59. Appender types (continued)

Type	Description
Rolling appender	Writes log statements to the file specified in the File Name field. Once the file size limit is reached (5 MB by default), the current file is renamed and a new file is created. For example, if the current file is named <code>maximo.log</code> , then the renamed file is <code>maximo.log.1</code> .
Daily Rolling appender	Writes log statements to the file specified in the File Name field. The file is renamed and a new file is created at a specified rate. This rate depends on the Date Pattern attribute. For example, if you have configured the Date Pattern of your Daily Rolling Appender to <code>yyyy-MM-dd</code> , when the current file is named <code>maximo_scheduled.log</code> , the renamed file is <code>maximo_scheduled.log.2007-06-18</code> .

Layouts

Layouts are components of the logging process. A layout determines the output format of a log statement.

A layout is always associated with an appender. For example, a Conversion Pattern such as: `%d{dd MMM yyyy HH:mm:ss:SSS} [%-2p] %m%n`, results in the following log statement: `2007-05-07 14:07:41,508 [main] INFO MyApp - Entering application;`

Loggers settings

Log settings define the type of information that is logged. Before you can activate logging for applications or runtime components, you must apply log settings in the Logging application.

If you create a logger or change the log settings, from the **Select Action** menu, select the **Apply Settings** action.

Perform this action if you change the settings in the **Log Level** field or in the **Active** check box. You also perform this action if you add loggers.

Log file locations

There are default locations for the log files to be stored in both IBM WebSphere Application Server and Oracle WebLogic Server.

If you use WebSphere Application Server, the default location of the log file is in the following folder:

```
\IBM\WebSphere\AppServer\profiles\ctgAppSrv01\maximo\logs
```

In this path, `ctgAppSrv01` is the profile name.

If you use WebLogic Server, the default location of the log file is in the following folder:

```
\BEA\92\user_projects\domains\base_domain\maximo\logs
```

In the second path example, `base_domain` can be the particular WebLogic Server domain that you configured.

You can also specify a new folder for storing your log files. However, if you set up a separate folder for your log files, you must ensure that the user account that you use has both read and write permissions on the folder.

Log file names

The corresponding log file name for an appender has a default value. However, when the file is created in the designated folder, the default file name is prefixed with the host name of the application server. The default file name is also prefixed with the server name of the system.

Example

The name of the log file of the Rolling appender is `maximo.log`. If the host name is `acme` and the server name of the system is `MXServer`, the file name is `acme_MXServer_maximo.log`. The server name is obtained from the value specified in the `maximo.properties` file that is part of the Enterprise Application Archive (EAR).

Loggers in multiple server environments

In a clustered environment, any logging changes that you make affect all servers in the environment. In a multiple server environment that is not clustered, these changes are applied only when you perform the **Apply Settings** action in the Logging application on each server.

If you use separate `logging.properties` files on separate servers in a clustered environment, rebuild and redeploy the `maximo.ear` file

EventTracker filter

The EventTracker filter is a Java class that can log all client events that are sent to the Tivoli's process automation engine. You can use the EventTracker filter to track the overall usage of application and identify the causes of potential issues.

After you enable the EventTracker filter, events are written to the `*_clientevents.log` in the `WAS appserver ROOT\profiles\profile\maximo\logs` directory, where

- `*` represents the combination of the hostname of the server and the Tivoli's process automation engine server.
- `WAS appserver ROOT` is the IBM WebSphere Application Server.
- `profile` is the profile under which the Tivoli's process automation engine was installed.

By default, the logging level for the EventTracker filter is set to `INFO`, which logs all events. To narrow the focus of logged events, you can set the logging level to `ERROR`, and then specify a specific user, application, or combination of user and application. If you set the logging level to `ERROR` and do not specify a user or application, then only error events are logged.

The EventTracker filter does not log exceptions. However, you can compare the timestamps in the error log to the EventTracker log to locate the potential causes of exceptions. Log entries are written in a tab-delimited format so you can import the log file into a spreadsheet application.

The EventTracker filter logs the following information:

Log item	Explanation	Example value
Time	The time that the event was received.	09/30/11 08:56:01.015
Duration	The time in milliseconds that was used to handle the event.	37
Server	The hostname combined with the IBM Maximo Asset Management server name.	localhost-MXServer
User	The user name of the user who sent the request.	wilson
MAXSessionID	The session ID for the user.	734
UISessionID	The user interface (UI) session ID for the request, which includes a number in parenthesis to indicate the total number of UI sessions the user has.	1(1)
App	The ID of the application that was used to send the request. The total number of applications in memory for a given UI session is listed in parenthesis next to the application ID.	wotrack(1)
Page	The ID of the page that the user had open when the event was sent.	mainrec
Sequence	The sequence value that was sent with the request. Multiple events can be sent in a single request, so if multiple logged events have the same sequence value, they were part of the same request.	2 1
Event	The client event that was sent to the server.	click
Target Control	The ID of the control that was the target of the event.	toolactions_ button_0
Target Component	The ID of the component that was the target of the event.	toolactions_ button_0- toolbarbutton_ image
Value	The value of the event.	wotrack
mxevent	When the click event is sent to the server, the mxevent provides additional information to explain the actual event that occurred. For example, when the user clicks on a Save button, the mxevent records that the event was sent to save the data.	INSERT
additionalevent	The additional event, if one was sent with the event.	
additionaleventvalue	The additional event value, if one was sent with the event.	

Working with logging

There are several tasks relating to managing your log files, such as specifying log file formats and locations, and associating loggers with appenders.

Creating logging.properties files

If you do not want to replicate the same logging configuration on all of your servers, maintain a separate logging.properties file on each server instance.

About this task

When you create a logging.properties file and deploy it in the application server of the system, you override the logging settings maintained in the database. In the Logging application, when the application server is running and you change log settings, these changes are only written back into the underlying database. The

next time you restart the application, the settings in the file (not the settings in the database) are applied. The settings in the file are applied because you have left a `logging.properties` file in the EAR.

Procedure

1. In the **Select Action** menu of the Logging application, select **List Logging Properties**.
2. Copy and paste the properties into any text editor available on the client computer.
3. Edit the properties in the text editor and save the contents as `logging.properties` file.
4. Copy the file into a system installation environment where you can generate a new Enterprise Application Archive (EAR) file to include the `logging.properties` file.
5. Copy the `logging.properties` file to the `maximo\applications\maximo\properties` folder.
6. After you generate a new EAR, deploy the EAR into the application server using standard EAR deployment steps.

Related concepts:

“Loggers” on page 347

Loggers are components of the logging process that prepare log statements to be written to console or log file.

“Loggers settings” on page 348

Log settings define the type of information that is logged. Before you can activate logging for applications or runtime components, you must apply log settings in the Logging application.

Specifying log file locations

You can specify a folder for storing your log files.

Before you begin

You must ensure that the user account used to run the application server has both read and write permissions on the log file destination folder.

Procedure

1. From the **Select Action** menu of the Logging application, select **Set Logging Root Folder**.
2. In the Set Logging Root Folder window, specify a location in the **Root Logging Folder** field.

Related concepts:

“Log file locations” on page 348

There are default locations for the log files to be stored in both IBM WebSphere Application Server and Oracle WebLogic Server.

Managing appenders

You can create, change, or delete appenders in the Logging application.

Procedure

1. From the **Select Action** menu of the Logging application, select **Manage Appenders**.

2. In the Manage Appenders window, you can create, modify, or delete an appender.

Automation scripts loggers

All logging that is related to scripts is done by the autoscript logger. The log level is set, by default, to ERROR. Only errors that are encountered during script execution, are output to the system console or to the product log file.

Individual script log levels can be used to redirect output statements in the script code to the system console or to the product log file. To ensure that output statements are redirected, set the log level of the individual script to the same log level of the autoscript logger. For example, if the autoscript logger is set to INFO, then the individual script log level must also be set to INFO.

The log level can be changed so that more script execution information is written to the product log file. Some of the statements are from the scripting component while others can be the script code-specific output statements. Each individual script can be configured at different log levels.

Example

A system administrator enters a purchase price value. A field validation script that is associated with the Purchase Price field of the Asset application runs business rules that compute the replacement cost. The system administrator encounters an error that indicates that the script failed to run line number 17 of the script. If required, he can place the output statements in the script code to isolate the specific section of code that is failing. He can set up detailed DEBUG logging to determine the type of variable, either input or output, that was received into or calculated in the script.

Cron task loggers

When cron tasks fail to run successfully, the root logger cron task can be set to different levels to obtain more detailed logs. If a particular cron task fails to run successfully, additional loggers can be set up to isolate the cause of the failure.

The following table lists the application area, the cron task, a description of the cron task, and the log levels to set each logger to. The list details the key cron tasks and not contain all of the cron tasks.

Table 60. Cron task loggers. Cron task logger details

Group / Area	Cron task	Description	Loggers and log level
Escalation	ESCALATION	Runs active escalation definitions and starts actions and notifications	To isolate execution errors with escalations, set the child logger of cron task root logger to DEBUG level.
Email listener	LSNRCRON	Runs periodically to process incoming emails and create or update Service Requests	To isolate execution errors with the email listener, set the child logger of cron task, EmailListnerCron to DEBUG level.

Table 60. Cron task loggers (continued). Cron task logger details

Group / Area	Cron task	Description	Loggers and log level
Integration	FLATFILECONSUMER	Automates the processing and loading of data contained in flat files	To isolate execution errors with the integration framework, set the integration root logger to DEBUG level. This logger supports log statements produced from all of the integration framework cron tasks.
Integration	IFACETABLECONSUMER	Automates the processing and loading of data contained in interface tables	To isolate execution errors with the integration framework, set the integration root logger to DEBUG level. This logger supports log statements produced from all of the integration framework cron tasks.
Integration	JMSQSEQCONSUMER	Automates the processing of messages from the JMS sequential queue in the integration framework	To isolate execution errors with the integration framework, set the integration root logger to DEBUG level. This logger supports log statements produced from all of the integration framework cron tasks.
Integration	XMLFILECONSUMER	Automates the processing and loading of data contained in XML formatted files	To isolate execution errors with the integration framework, set the integration root logger to DEBUG level. This logger supports log statements produced from all of the integration framework cron tasks.
Reporting	REPORTLOCKRELEASE	Determines whether locked report jobs continue, or are canceled and resubmitted to the reporting queue	To isolate execution errors with reporting, set the child logger REPORTLOCKRELEASE for the cron task root logger to DEBUG.
Reporting	REPORTOUTPUTCLEANUP	Automates the clean-up of previously output report entries that are stored	To isolate execution errors with reporting, set the child logger REPORTOUTPUTCLEANUP for the cron task root logger to DEBUG.
Reporting	REPORTSCHEDULE	Automates the removal of report usage entries that are created each time a report is run	To isolate execution errors with reporting, set the child logger REPORTSCHEDULE for the cron task root logger to DEBUG.
KPI	KPICronTask	Refreshes KPI values periodically	To isolate execution errors with KPI, set the child logger KPICronTask for the cron task root logger to INFO.
User group/ synchronization	LDAPSYNC	Automates the synchronization of users and groups maintained in an LDAP directory	To isolate execution errors with security, set the child logger LDAPSYNC for the cron task root logger to DEBUG.

Table 60. Cron task loggers (continued). Cron task logger details

Group / Area	Cron task	Description	Loggers and log level
User group/synchronization	VMMSYNC	Automates the synchronization of users and groups maintained in an LDAP directory	To isolate execution errors with security, set the child logger VMMSYNC for the cron task root logger to DEBUG.
Inventory	ReorderCronTask	Runs the reorder process for direct issue items	To isolate execution errors with inventory, set the child logger, INVENTOR for the application root logger to DEBUG.

Related information:

 [Key Report Property Settings](#)

Escalation loggers

If an escalation fails to run, the failure can be caused by the escalation engine, the email notifications, or actions triggered by the escalation. You can troubleshoot the cause of the failure by setting the correct log level.

The following table lists the escalation loggers, a description of the logger, and the log level to set each logger to.

Logger	Description	Log level
crontask	Root logger that generates log statements for all cron tasks	DEBUG
ESCALATION	Child logger that inherits from cron task root logger and generates log statements specifically for the escalation engine	DEBUG
sql	Root logger that generates log statements of SQL statements run in the application server	INFO
COMMTEMPLATE	Child logger that inherits from service root logger and generates log statements specifically for communication templates and notifications events	INFO
mail	Root logger that generates log statements for communication with the mail server when sending notifications.	DEBUG

Integration framework loggers

The integration framework has four root loggers: integration, REST, interaction, and Open Services Lifecycle Collaboration (OSLC). To receive the maximum amount of information set the root loggers to DEBUG.

The following table lists the loggers, a description of the log level, and log levels to set the logger to.

Table 61. Integration framework loggers. Integration framework logger details

Logger	Description	Log level
Integration	Logs information about the integration framework configuration, creation, management, and runtime processing.	DEBUG
REST	Logs information about the REST API components. REST API starts the integration with business objects. The REST API call from a client, runs the client request against the appropriate business object and returns a response.	DEBUG
Interaction	Logs information about both the design process, when a web service interaction is created, and the execution process, where a web service interaction is run. Designated users run web service interactions to receive data interactively from external systems."	DEBUG
OSLC	Logs information about both the design process, when an OSLC interaction is created, and the execution process, where an OSLC interaction is run.	DEBUG

Enabling the EventTracker filter

To enable the EventTracker filter, uncomment the filter definition and filter mapping in the web.xml file, which is located in *Tpae root*\applications\maximo\maximouiweb\webmodule\WEB-INF\ directory.

About this task

If your web.xml file does not contain the commented code, you can enable the EventTracker filter by adding the code, without the comments. Add the filter definition between the last filter definition and the first filter mapping in your web.xml file. Then, add the filter mapping so that it is the last entry in the filter mappings.

Procedure

1. Open the web.xml file from the *Tpae root*\applications\maximo\maximouiweb\webmodule\WEB-INF\ directory.
2. Remove the comment from the following code for the filter:

```

<!-- Uncomment these lines to enable the EventTracking filter
<filter>
  <filter-name>EventTrackingFilter</filter-name>
  <filter-class>psdi.webclient.system.filter.EventTrackingFilter</filter-class>
</filter>
-->

```

3. Remove the comment from the following code for the filter mapping:

```

<!--! Uncomment these lines to enable the EventTracking filter
<filter-mapping>
  <filter-name>EventTrackingFilter</filter-name>
  <url-pattern>/ui/*</url-pattern>
</filter-mapping>
<filter-mapping>
  <filter-name>>EventTrackingFilter</filter-name>
  <url-pattern>/webclient/login/logout.jsp</url-pattern>
</filter-mapping>
-->

```

4. Save and rebuild the Tivoli's process automation engineenterprise archive (EAR) file.
5. Redeploy the EAR file.

What to do next

In the Logging application, a new row in the Root Loggers table lists the EventTracker filter as **eventtracking** with the key value **log4j.logger.maximo.webclient.eventtracking**. By default, the logging level is set to INFO, which logs all events. If the logging level is set to ERROR, change the logging level to INFO and select the **Apply Settings** action.

Logging events for specific applications or users

You can add rows to the Logger table in the Logging application to enable the EventTracker filter to log the events of specific applications, specific users, or specific users in specific applications.

Procedure

1. In the Logging application, in the Root Logging table, set the logging level for **log4j.logger.maximo.webclient.eventtracking** to ERROR.
2. In the Logging table, add a new row.
3. Set the **Logger** field to one of the following:

Option	Description
To log a specific application	Specify the application ID, for example, wotrack.
To log a specific user	Specify the user ID, for example, wilson.
To log a specific application for a specific user	Specify the application ID.user ID, for example, wotrack.wilson.

The **Key** field is automatically populated based on the value in the **Logger** field.

4. From the **Select Action** menu, select the **Apply Settings** action.

Stopping the logging of events

You can stop logging events in the Logging application, which eliminates the need to edit the web.xml file and restart the application server.

Procedure

1. In the Logging application, in the Root Logging table, set the logging level for **log4j.logger.maximo.webclient.eventtracking** to ERROR.

2. From the **Select Action** menu, select the **Apply Settings** action.

Log correlation

Correlation is the process of analyzing a set of related events. The analysis is based on rules that are used to interpret the event data. When log correlation is enabled, you can use the correlation ID to identify the UI request, cron task action, MXSCRIPT object, or REST call that produced a log entry.

Correlation ID

When log correlation is enabled, identify the user interface (UI) request, cron task action, MXSCRIPT, or REST call that produced a log entry by using the correlation ID. All log entries for a single UI request or cron task action have the same correlation ID. The unique correlation ID applies to any group of log entries such as SQL logs or application logs.

The correlation ID is in the format of *[CID-TYPE-NUMBER]*. The type can be CRON, UI, MXSCRIPT, or REST. The number ensures uniqueness. The server name is also included when log correlation is enabled. For log entries that are not yet correlated, the correlation ID field in the log is empty. The server name field is empty for logs that are generated before the server name is available.

To enable log correlation, in the System Properties application, add the **mxe.logging.CorrelationEnabled** property.

To enable UI request correlation, in the System Properties application, add the **mxe.webclient.logging.CorrelationEnabled** property.

Example

The following example shows the correlation log for the UI requests that are associated with an SQL statement that is being debugged:

Interpret the log file

```
06 Jun 2012 16:12:38:360 [INFO] [MXServer] [CID-UI-993] Correlation started.
06 Jun 2012 16:12:38:534 [INFO] [MXServer] [CID-UI-993] BMXAA6719I - USER = (WILSON) SPID =
(137) app (WOTRACK) object (WORKORDER) : select * from workorder where (woclass in (select
value from synonymdomain where domainid = 'WOCLASS' and maxvalue in ('WORKORDER','WOACTIVITY'))
and historyflag = 0 and istask = 0 and siteid = 'BEDFORD')
06 Jun 2012 16:12:38:546 [INFO] [MXServer] [CID-UI-993] BMXAA6719I - USER = (WILSON) SPID =
(149) app (WOTRACK) object (WORKORDER) : select * from workorder where (workorderid = 202)
06 Jun 2012 16:12:38:582 [INFO] [MXServer] [CID-UI-993] BMXAA6719I - USER = (WILSON) SPID =
(27) app (WOTRACK) object (WORKORDER) : select * from workorder where (workorderid = 202)
06 Jun 2012 16:12:38:717 [INFO] [MXServer] [CID-UI-993] BMXAA6719I - USER = (WILSON) SPID =
(20) app (WOTRACK) object (WORKORDER) : select count(*) from workorder where
(woclass in (select value from synonymdomain where domainid = 'WOCLASS
' and maxvalue in ('WORKORDER','WOACTIVITY')) and historyflag = 0 and istask = 0 and siteid
= 'BEDFORD')
06 Jun 2012 16:12:38:998 [INFO] [MXServer] [CID-UI-993] Correlated data: BEGIN UIsessionId:3
Event:loadapp AppName:autoscript UserId:wilson UIClientIP:127.0.0.1 ElapsedTime:638 ms END

06 Jun 2012 16:12:44:589 [INFO] [MXServer] [CID-UI-994] Correlation started.
06 Jun 2012 16:12:44:597 [INFO] [MXServer] [CID-UIASYNC-995] Correlation started.
06 Jun 2012 16:12:44:634 [INFO] [MXServer] [CID-UIASYNC-995] BMXAA6719I - USER = (WILSON)
SPID = (149) app (WOTRACK) object (WORKORDER) : select * from workorder where ((woclass in
(select value from synonymdomain where domainid = 'WOCLASS' and maxvalue in ('WORKORDER',
'WOACTIVITY')) and historyflag = 0 and istask = 0 and siteid = 'BEDFORD')) and ((status like
'%WAPPR%'))
06 Jun 2012 16:12:44:641 [INFO] [MXServer] [CID-UIASYNC-995] Correlated data:
BEGIN UIsessionId:3 AppName:wotrack UserId:wilson UIClientIP:127.0.0.1 ElapsedTime:44 ms END
```

- The first log statement in the fragment displays a field **[CID-UI-993]** that is the component that is correlated in this example.

- CID is the correlation ID, followed by the component that is being correlated, and the numeric value that is applied to all the correlated statements output for the particular correlation.
- The [CID-UI-993] field groups a set of correlated log statements that are started by a UI request.
- A correlated set of log statements are marked by correlation started and correlation data statements. The first set of correlated log statements begins with [CID-UI-993]correlation started and ends with [CID-UI-993]Correlated data:BEGIN the specific component information END.
- The second set of correlate log statements indicate that the next set of correlated log statements is being output.
- When a UI request is correlated, sql commands that run in response to user interaction are placed in the same group of correlated log statements as the UI request.
- In the second correlated statement for correlation [CID-UI-993], the sql command is associated with the WORKORDER table.
- If multiple sql statements run in response to the UI request the statements are output to the same group of correlated log statements.
- The final log statement for correlation [CID-UI-993] begins with Correlation data and provides information about the user session that started the output of correlated log statement.
- The additional information includes:
 1. The UI session Identifier 3.
 2. The UI event loadapp.
 3. The application that started the UI event autoscript.
 4. The user that started the event wilson.
 5. The IP address from where the UI request was received UIClientIP.
- In the final log statement for correlation [CID-UI-993], the ElapsedTime information specifies the time spent by the application server to process the UI request.

Chapter 15. Managing bulletin boards

You use bulletin boards to post messages about critical problems and incidents, and to broadcast information throughout an organization. You can also view communication logs from service desk agents and their users, and communications that are generated by workflow processes and escalations.

Bulletin board overview

You can create messages about critical problems and incidents that you can broadcast throughout an organization. Recipients can view bulletin board messages from the navigation bar of an application or from the Start Center.

Communication logs for bulletin board messages

You can view a communication log for bulletin board message in the Bulletin Board application. Communication logs are available in several applications.

The log contains communications about outbound messages that were sent between service desk users and agents. You also view communications that were generated by escalations and workflows if the related communication template specifies that these messages are stored in logs.

Working with bulletin boards

Using the Bulletin Board application, you can create bulletin board messages, specify the audience for the messages, and perform other tasks.

Viewing bulletin board messages

You can use the bulletin board to create, view, and post messages related to critical problems and incidents, and to broadcast company-wide information. Creating and posting messages on the bulletin board minimizes the creation and duplication of tickets.

Before you begin

You must be granted access to the Bulletin Board application before you can create and post messages.

About this task

Messages can be routed to all users, or just selected organizations, sites, or person groups. You can view only those bulletin board messages for which you have been selected as an audience. If the composer of the message has not selected a specific audience, any user who logs in to the system can view the message.

Located in the Navigation bar of all applications, the icon named **Bulletins** indicates whether you have unread bulletin board messages. The icon uses a numeric value to indicate the number of bulletin board messages that you can view. The icon also changes from red to green when a new message is posted.

Procedure

1. In the Navigation bar, the Bulletin Board area displays all the bulletin board messages. When you change the value in the **Viewed** field, you change which messages you see. By default, the value in the **Viewed** field is **N**, and the Bulletin Board area displays unread messages. To view either previously viewed messages or all messages:
 - Enter **Y** in the **Viewed** field to see previously viewed Bulletin Board messages.
 - Leave the **Viewed** field blank to see all messages (both previously viewed and unread bulletin board messages). The messages display, with the most recent message appearing at the top of the list. Each message displays the date and time when the message was posted to the Bulletin Board.
2. In the Bulletin Board window, click a message to display the details of the message.
3. Click the message again to collapse the detailed view of the message.
4. Click **OK**.

Creating bulletin board messages

You can create messages to communicate to users about critical problems and incidents, and to broadcast information throughout the organization.

Procedure

1. In the Bulletin Board application, click **New Message**.
2. Type an identifier, a subject for the message, a posting date, and an expiration date.
3. Type a message of up to 4,000 characters.
4. Optional: Limit the message to an audience by organization, by site, or by person group.
5. Save your changes.

Specifying audiences for bulletin board messages

You can specify the audience that can view a bulletin board message. You can limit the audience by organization, by site, or by person group. If you do not specify an audience, any user can view the message.

Procedure

1. In the Bulletin Board application, open the message for which you want to specify an audience.
2. Click a tab to specify an audience by organization, by site, or by person group.
3. On the tab, select one or more organizations, sites, or person groups and save your changes.

Changing the status of bulletin board messages

You can change the status of a bulletin board message to control the content of what is posted to the bulletin board. The status of a bulletin board message can be draft, approved, or rejected.

Procedure

1. In the Bulletin Board application, open the message for which you want to change the status.

2. From the **Select Action** menu, select **Change Status**.
3. In the Change Status window, select a status and click **OK** to save your changes.

Example

For example, you can create a bulletin board message with a status of draft, and then route the message to a supervisor. The supervisor reviews the message, and changes the status to approved. The message is then posted to the bulletin board.

Copying bulletin board messages

You can copy an existing bulletin board message. Once you copy a message, you can use the same information that is in the original message. You can also change the values in the message that you copied and save it as a new message.

Procedure

1. In the Bulletin Board application, open the message that you want to copy.
2. From the **Select Action** menu, select **Duplicate Message**.
3. Optional: In the new message, change the subject, the message, and the posting and expiration dates.
4. Save your changes.

Viewing communication logs for bulletin board messages

You can view outbound bulletin board messages and any attachments that were sent between users and service desk agents.

Procedure

1. In the Bulletin Board application, open the message for which you want to view the communication log.
2. Click the **Communication Log** tab. By default, you can view the messages in reverse chronological order according to the **Date** field.
3. Click **View Details** to view the details of a message. The information in the communication log is read-only.

Related concepts:

“Communication logs for bulletin board messages” on page 359

You can view a communication log for bulletin board message in the Bulletin Board application. Communication logs are available in several applications.

Viewing the history of bulletin board messages

Each time you change the status of a bulletin board message, the change is stored in a history transaction. You can view this read-only status history of a bulletin board message.

Procedure

1. In the Bulletin Board application, display the message whose status history you want to view.
2. From the **Select Action** menu, select **View History**.
3. View the status information and click **OK**.

Deleting expired bulletin board messages

By default, expired bulletin board messages are not deleted. If you do not need to keep expired messages, you can change your settings so when expired the messages are deleted.

About this task

To activate this setting, you must change the `mxe.crontask.deleteBB` system property and add parameters for the ESCBLTNEXP escalation.

Procedure

1. In the System Properties application, open the `mxe.crontask.deleteBB` system property.
2. Specify the global property value as 1, save the property, and refresh the session.
3. In the Escalations application, open the ESCBLTNEXP escalation.
4. In the Select Action menu, click **Activate/Deactivate Escalation** to deactivate the escalation.
5. Specify values in the **Elapsed Time Interval** field and in the **Interval Unit of Measure** field for the EXPIREDATE elapsed time attribute.
6. In the Actions section, add the DELEXPMSG action.
7. Save the escalation.
8. Activate the escalation.

Chapter 16. Working with sets

You use sets to create item sets and company sets that are shared across multiple organizations.

About this task

The following relationships apply to sets and organizations:

- You must create at least one item set and one company set before you can create any organization.
- You can create as many item sets and company sets as your business practices require.
- You must associate each organization with only one company set.
- You must associate each organization with only one item set.

You cannot delete an item or company set if the set is not empty or if the set is associated with an organization.

Creating item sets or company sets

You create an item or company set to enable the sharing of item master and company master information between organizations.

Before you begin

In the Sets application, you view organizations that are associated with sets. Before you can view the organization, you must use the Organizations application to assign the set to at least one organization.

About this task

You can create an unlimited number of sets. Multiple organizations can use the same item set or company set.

Procedure

1. In the Sets application, click **New Row**.
2. Provide the name of the set. The name must be unique for all sets; an item set and a company set cannot have the same name.
3. Optional: Provide a description of the set.
4. Specify a value for the type.
5. Optional: If you specified COMPANY as the value for type, select the **Automatically Add Companies to Company Master** check box to automatically create a company master record. A company master record is then created whenever a user adds a company in the Companies application. If you do not select this check box, users must add companies in the Company Master application.
6. In the **Default Item Status** field, specify the status that you want new item master, service item, and tool records to have.

7. Click **Save Sets**. After you save the record, the only fields that you can edit are the **Sets Description** field and, for a company set, the **Automatically Add Companies to Company Master** check box.

What to do next

After you create a set, you use the Organizations application to assign the set to an organization.

Changing item or company sets

After you create an item or company set, you can change settings to suit your business needs.

Procedure

1. In the Sets application, click **View Details** for the set that you want to change.
2. Change one or more of the following settings:
 - Change the description.
 - For a company set, change the **Automatically Add Companies to Company Master** check box.
3. Save your changes.

Chapter 17. Managing organizations

You can set up the organizations and sites that you use. You can set defaults for various options that are related to applications or groups of applications, such as Work Order Tracking, Inventory, Purchasing, Preventative Maintenance, Assets, and so on.

About this task

The following relationships apply to sets and organizations:

- You must create at least one item set and one company set before you can create an organization.
- You can create as many item sets and company sets as your business practices require.
- You must associate each organization with only one company set.
- You must associate each organization with only one item set.

Organizations overview

You can set up multiple organizations and sites for your company. The divisions are determined by the types of operations that are performed at different locations, and what data can be shared among them. Different organizations and sites keep parts of their operations separate, while sharing others.

Application levels and data storage

Application data is stored at different levels. The separation of data allows you to reuse as much data as possible across organizations and sites without breaching best practices in security and authorization.

The following are the four application levels:

Table 62. Application levels

Application levels	Description
System level	The data is available to all organizations and sites.
Set level	This level is a special category by which multiple organizations can share items and vendor company data. The system stores this data by default at the organization level.
Organization level	The data is available to only the specified organization and all sites within the organization.
Site level	The data is available to only the specified site.

The following examples illustrate how the level at which application data is stored affects how you use the system:

- The application level determines the level at which record IDs must be unique.

For example, Work Order Tracking is a site-level application; two different sites can both use the same work order number to identify two different work orders. Chart of Accounts is an organization-level application; two different organizations can use the same general ledger account number to specify different general ledger accounts.

- The application level affects some aspects of security.

For example, if a security group has access to one site for an organization-level application, then members of that group can access all of the application records for that organization.

- The application level affects some user default settings.

For example, if a site is used as a filter to display records, but the application is at the organization level, then a user can access all records for the organization that owns that site.

Sites and organizations

A site is a division within an organization that maintains certain data independently from other sites. When you create an organization, you create at least one site.

When you create a site, the values that you specify are used as defaults for that site. You can change the values as needed. You cannot delete a site and you cannot delete an organization after you add a site.

You can use sites to administer security, and to give users different rights at different sites.

Sites that belong to the same organization must use the same currency. They must also share the same options for work orders, assets, labor, and other types of data.

Activation and deactivation of organizations and sites

The ability to activate and deactivate organizations and sites is dependent on the enterprise hierarchy. When you create an organization, you create at least one site. When you deactivate an organization, all sites within that organization become inactive. This ensures that the structure of your company is not compromised.

When you deactivate an organization or site, users can no longer create records for that organization or site. A deactivated organization or site and all its information is not deleted. Records that reference a deactivated organization or site are not affected.

If a site and its organization are both deactivated and you activate the organization, the site is not automatically activated.

You can activate and deactivate an organization or a site on the **Organization** tab of the Organizations application.

Item sets

Item sets let multiple organizations within a company view and choose from a common set of items stocked in the storeroom. An organization can be associated with only one item set. However, multiple organizations can use the same item set.

You identify an item set with a unique name or number, and each item that belongs to the item set has a unique item number.

When you create an organization, you associate an item set to it. When you create an item, by default the item is placed into the item set associated with the organization for your default insert site.

Autonumbering

When you set up autonumbering and a user creates a record, the record IDs or other specified fields increment by one.

Many applications, such as Work Order Tracking, Purchase Requisition, Purchase Orders, and Invoices have autonumbering set up for their key fields such as record ID. You can change the default starting sequence. You can also add a prefix to an existing numbering sequence or add autonumbering to other fields. For applications that do not have autonumbering, you can set up autonumbering for the record ID field or other fields.

Applications store data at one of these levels: system, set, organization, and site. You use different actions to set up autonumbering for applications at each level. For a multisite implementation, you can have different autonumbering sequences for different organizations and sites. For example, Work Order Tracking is a site-level application. You can set up different autonumbering sequences for each site.

You specify the numbering scheme for autonumbering in the Organizations application. To implement the numbering scheme in the organization that you are autonumbering, use the Database Configuration application.

System-level autonumbering

You can specify autonumber seeds and prefixes for applications at the system level.

For example, Incidents is a system-level application. Incident IDs are unique at the system level. If you implement autonumbering for incident records with a seed of 1000, each new incident record increments by one. This increment occurs regardless of where the incident record is entered. If incident 1000 is inserted at site 1 in organization A, and the next incident record is inserted at site 3 in organization B, then the incident number at site 3 will be 1001.

Organization-level autonumbering

If you want all organizations to be part of the same numbering sequence, you can use the Organizations application to implement autonumber seeds and prefixes for organization-level applications.

For example, Calendars is an organization-level application. Calendar IDs are unique at the organization level. If you specify autonumbering for calendars in organization A with a seed of 1000, then each new calendar record increases by one in the organization, regardless of which site generates the calendar. If calendar 1000 is inserted at site 1, and the next calendar record is inserted at site 2, the calendar number at site 2 is 1001.

Site-level autonumbering

If you want all sites in the organization to be part of the same numbering sequence, you use the Organizations application to implement autonumber seeds and prefixes for site-level applications.

For example, Preventive Maintenance is a site-level application and preventive maintenance IDs are unique at the site level. However, you can implement an organization-level sequence. You can choose to have the autonumbering work in the following ways:

- All sites are part of the same sequence. You set the autonumbering for preventive maintenance records and the autonumbering works the same as with an organization-level application. If preventive maintenance 1000 is inserted at site 1, and the next preventive maintenance record is inserted at site 2, the preventive maintenance number at site 2 is 1001.
- You have three sites and you want one to have independent preventive maintenance autonumbering. You want sites 1 and 2 to be on the same autonumber sequence, but you want site 3 to be independent. You set preventive maintenance autonumbering to 1000. You then specify an independent autonumber sequence for site 3.
- You want all sites to have independent preventive maintenance autonumbering. You specify autonumbering for all sites, or you specify autonumbering for all sites except one. The exception is defined by whatever you specify in the Organization Level window.

Set-level autonumbering

In the Organizations application, you can specify autonumber seeds and prefixes for applications at the set level.

For example, Item Master is a set-level application, and item master IDs are unique at the set level. You implement autonumbering for item master records with a seed of 1000. Each new item master record increments by one, regardless of where the item master record is added. If item master 1000 is inserted at site 1 in organization A, and the next item master record is inserted at site 3 in organization B, then the item master number at site 3 is 1001.

ABC breakpoints and organizations

ABC analysis helps you quickly identify which inventory items represent the greatest monetary investment for your company in terms of dollar value and rate of turnover. This analysis ensures that crucial inventory does not fall below minimum levels, and helps keep current balance figures for an item reconciled with the actual count.

The ABC type value for an item is determined by running the Inventory ABC Analysis report. You can run this report from the Inventory application or module. This report multiplies the current year-to-date issued quantity by the last cost of the item. The report sorts the items in descending order of the calculated dollar value.

The types values for A, B, or C are set based on the percentages that you specify in the Organizations application. You can set an ABC type to indicate no category, and that type is excluded from the ABC analysis.

For example, you can use these parameters:

- Type A break point 30% (enter .30)
- Type B break point 30% (enter .30)
- Type C break point 40% (enter .40)

Using these values, the top 30% of items by cost is type A, the next 30% of items is type B, and the last 40% is type C.

The Inventory ABC Analysis report also changes the cycle count for the item, based on the ABC type and the cycle count. If the ABC type is set to N, the cycle count is not changed.

You set the cycle count to the number of days for each ABC type. For example, type A cycle counts can be 30 days, type B cycle counts 60 days, and type C cycle counts 90 days.

You can run the Inventory Cycle Count report to show each item that is due for a cycle count in the next month, together with its ABC type.

Enablement of repair facilities

To create work orders in a site that is different from the site of the asset on the work order, enable the repair facilities feature.

Enablement of repair facilities is useful for transportation assets and other moveable assets. Moveable assets are often serviced on the road, at a repair facility that does not belong to the home site of the asset.

Enablement of the feature means that a repair facility, a special type of location, can take ownership of work orders from multiple sites that are in the same organization. User security can then be configured to give permissions to view work orders in multiple sites if the work orders are owned by the repair facility. When you enable the feature, fields that are related to repair facilities are made visible in the Work Order Tracking application and other affected applications.

The **Enable Repair Facilities** check box is on the Other Organization Options window in the Organizations application.

If you enable the repair facilities feature, the **Bypass Site Mismatch Warning Message** check box on the Other Organization Options window becomes read/write. Select the check box to prevent the display of a warning message when a work order is created in a site that is different from the site of the asset on the work order.

Customization options for applications

There are many options and settings that you can customize for your applications in the Organizations application. They include work order options, inventory, purchasing, asset, and preventive maintenance (PM) options. You can access the options from the Select Action menu.

Table 63. Customization options

Options	Description
Work orders	Defines and changes work types and work order types for classes of work orders.
Edit rules	Specifies the work order properties that can be changed for a given work order status. For example, you specify that an asset and a location can be edited when it is approved.
Actual start date for work orders	Specifies whether the actual start date is the date when the work order status is changed to initiate or changed to complete.

Table 63. Customization options (continued)

Options	Description
Default problem start time	Specifies whether the downtime starts when someone reports the problem or when the problem occurs. The default date and time are shown in the Start field in the Downtime Report window.
Site options	Specifies how to increment task numbering such as the starting number for the first task and the increment value for each task. Task numbering is used on the Plans tab in the Work Order Tracking application
Inventory defaults	Specifies default settings for inventory such as ABC breakpoints, cost and currency variances, and negative current and available balances.
Inventory reorders	Specifies whether an approved or unapproved purchase requisition or purchase order is created when a reorder request is generated. You can also specify the maximum number of reorder lines on a purchase order or purchase requisition.
Inventory costs	Specifies the default issue cost that applies when an item is issued from a selected site.
Transfer options	Specify whether a shipment record is required to transfer items.
Purchasing defaults	Specify default settings for the purchasing life cycle such as whether a PR must be approved before it is converted to a PO or contract. You can also specify if costs can be added to or modified for line items.
Purchase order labor	Specify whether a purchase order is required to approve internal labor.
Tax options	Specify the default tax general ledger accounts and tax codes to calculate the amount of tax that is due on a purchase requisition, request for quotation, purchase order, or invoice at the organization level. You can also specify the order in which tax codes are used.
Contract options	Specify contract options such as contract types, and the terms and properties that are associated with contract types.
Invoice options	Specify default settings for invoices such as whether to validate the financial period when you save an invoice. You can also copy the invoice date specified by the vendor to the General Ledger (G/L) posting date.
Asset options	Specify whether to record asset history when the status of a work order changes to complete or to closed.
Drilldown options	Specify how a drilldown opens for locations and assets. Specify that the drilldown top level starts at the top-level location or asset.
Preventive maintenance (PM) options	Specify general preventive maintenance (PM) options such as whether the priority entered on job plans is used to sequence PMs. You can also specify whether the frequency criteria on a PM is used to generate work orders. You can also provide the lead time in days between when an automatically generated work order is created and when the preventive maintenance is scheduled.
Safety plan options	Specify whether data for hazards is available with the work assets data in the Safety Plans application.

Table 63. Customization options (continued)

Options	Description
Labor options	Specify default settings for labor transactions such as the approval process for internal and external labor. You can also specify whether craft and labor can be different or must match when you report actuals.
Global ticket solutions options	Specify whether solutions for global tickets are automatically applied to related global tickets.
Service level agreement options	Specify how you want to match service level agreements (SLAs) to records for example, one or multiple SLAs. You can also specify the target dates to apply to a record.
Workflow options	Specify workflow processes for work orders, purchase requisitions, and purchase orders.
System settings	Specify general system settings such as the default start and end dates, and the format of names. You can also specify the character that represents the unspecified general ledger components in an account code.

Taxes for organizations

You can specify default tax general ledger accounts and tax codes to calculate the amount of tax that is due on a purchase requisition, request for quotation, purchase order, or invoice. You can also specify the order in which tax codes are used in calculations.

A tax type can be, for example, a federal, state, or city sales tax. Another tax type might be a special tax for handling hazardous material.

A tax code represents a particular tax rate, such as MA for the Massachusetts sales tax of 5%. Therefore, one tax type might include tax codes for all state or provincial sales taxes.

The requirements of your financial system determine how you specify the tax types, and also how many tax types you use. In the Tax Options window of the Organizations application, the default number of tax types that you can view is five. If your financial requirements dictate additional tax types, you can use the Database Configuration application to specify up to 27 different tax types. Each tax type can have any number of tax codes.

If your financial system uses one, you can specify a general ledger account for paid and unpaid taxes.

Taxes paid

Taxes paid to vendors

Unpaid taxes

Taxes not yet paid to the government

Adding taxes to the cost of an item increases the stocked item average cost for that item. To avoid this increase, you can add the tax to issue on receipt items only. For items issued directly to a work order, general ledger account, location, or asset, the tax is only added to the cost of the individual item.

Drilldown options

Drilldown options specify how a drilldown opens. The drilldown options apply to applications in which both the **Location** field and the **Asset** field on a record are blank. The options also apply to the **Open Drilldown** action in the Assets application.

Drilldown includes two options:

Top Level Starts at Top Level Location (Start at Top Location in Primary System)

The drilldown opens to the **Location** tab, and begins with the top-level location of the primary system. This option is useful for organizations with many asset records that use location hierarchies to organize these records.

Top Level Starts at Top Level Asset (Show All Assets without Parent)

The drilldown opens to the **Asset** tab. This option is useful when asset records are organized into a hierarchy with a few top-level asset records. If the organization has many top-level asset records, selecting this option might slow or hinder performance when you display the records.

Working with organizations

You can set up the organizations and sites that you use. You can set defaults for various options that are related to applications or groups of applications, such as Work Order Tracking, Inventory, Purchasing, Preventative Maintenance, Assets, and so on.

Creating organizations

In the Organizations application, you create the organizations to use with the system. Each organization can have multiple sites.

Before you begin

You must create at least one organization and one site. Before you create the first organization, use the Currency Codes application to create at least one currency code. You cannot change the base currency after you save the record. Use the Sets application to create at least one item set and one company set. Use the Chart of Accounts application to create a general ledger account. Multiple organizations can use the same item set or company set.

About this task

Multiple organizations and sites can use the same database, while certain business procedures and information remain specific to the individual organizations and sites. After you save an organization record, you cannot edit the item or company set. After you add a site to the organization and save the record, you cannot edit base currency 1.

Procedure

1. In the Organizations application, click **New Organization**.
2. Type a name for the organization.
3. Specify base currency code 1 for your base currency.
4. Optional: Specify base currency 2.
5. Specify the item set and the company set that you want to associate with this organization.

6. In the **Default Item Status** field, select the status that you want new item master, service item, and tool records to have.
7. Click **Save Organization**.

Activating organizations

When you create an organization, it is inactive by default. To allow users to interact with an organization, you must activate it.

Before you begin

To activate an organization, you must first create the account in the Chart of Accounts application. This process includes creating general ledger component values and general ledger codes, and activating general ledger components.

About this task

You can activate any sites that are inactive.

Procedure

1. In the Organizations application, select the organization that you want to activate.
2. On the **Organization** tab, in the **Clearing Account** field, specify the account code that you created in the Chart of Accounts application.
3. Select the **Active** check box.

Deleting organizations

As your business needs change, you can delete the organizations that you created.

About this task

You can delete an organization only if it has no sites associated with it.

Procedure

1. On the **List** tab in the Organizations application, select the relevant organization.
2. From the **Select Action** menu, select **Delete Organization**.
3. Save your changes.

Clearing material reservations for work orders

To customize organization properties specific for your business needs, you can specify when to clear the material reservation for a work order.

Procedure

1. In the Organizations application, select an organization.
2. From the **Select Action** menu of the **Organization** tab, click **Work Order Options > Other Organization Options**.
3. In the Other Organization Options window, specify whether to clear the material reservation when the work order status changes to either Complete or to Closed.
4. Click **OK**.

Specifying options for work orders and ticket owners

In the Organizations application, you can determine the behavior of the **Select Owner** action in the Work Order Tracking application and in several service desk applications. You can specify whether there is a value in the **Date** field in these applications, and if the list of persons is filtered according to availability.

Procedure

1. In the Organizations application, select the organization for which you want to specify ownership options.
2. From the **Select Action** menu, select **Ownership Assignment Options**.
3. In the Ownership Assignment Options window, select a site.
4. Select one of the following options.

Option	Description
Do not check Person Availability	In the Work Order Tracking application and relevant service desk applications, the Date field is blank and the list of persons is not filtered by availability.
Check Person Availability	In the Work Order Tracking application and relevant service desk applications, the Date field is populated. The list of persons is filtered by availability at that date, according to calendars.

5. Click **OK**.

Related reference:

“Customization options for applications” on page 369

There are many options and settings that you can customize for your applications in the Organizations application. They include work order options, inventory, purchasing, asset, and preventive maintenance (PM) options. You can access the options from the Select Action menu.

Setting purchasing options

You can customize purchasing options for your specific business needs.

Related concepts:

“Taxes for organizations” on page 371

You can specify default tax general ledger accounts and tax codes to calculate the amount of tax that is due on a purchase requisition, request for quotation, purchase order, or invoice. You can also specify the order in which tax codes are used in calculations.

Related reference:

“Customization options for applications” on page 369

There are many options and settings that you can customize for your applications in the Organizations application. They include work order options, inventory, purchasing, asset, and preventive maintenance (PM) options. You can access the options from the Select Action menu.

Associating properties with contracts for organizations

You can associate properties with a contract.

Procedure

1. In the Organizations application, select the relevant organization.

2. From the **Select Action** menu of the **Organization** tab, click **Purchasing Options > Contract Options**.
3. In the Contract Options window, select the contract type.
4. Click **Associate Properties**.
5. In the Associate Properties window, select a property to associate with the contract type.
6. Optional: Specify a default value for the property and specify whether the property can be edited.
7. Click **OK**.

Related concepts:

“Taxes for organizations” on page 371

You can specify default tax general ledger accounts and tax codes to calculate the amount of tax that is due on a purchase requisition, request for quotation, purchase order, or invoice. You can also specify the order in which tax codes are used in calculations.

Related reference:

“Customization options for applications” on page 369

There are many options and settings that you can customize for your applications in the Organizations application. They include work order options, inventory, purchasing, asset, and preventive maintenance (PM) options. You can access the options from the Select Action menu.

Associating terms and conditions with contracts for organizations

You can associate terms and conditions with a contract.

Before you begin

The terms and conditions that you associate are derived from values in the Terms and Conditions application.

Procedure

1. In the Organizations application, select the relevant organization.
2. From the **Select Action** menu of the **Organization** tab, click **Purchasing Options > Contract Options**.
3. In the Contract Options window, select the contract type.
4. Click **Associate Terms**.
5. In the Associate Terms and Conditions window, click **New Row**.
6. In the **Terms** field, specify a value.
7. Click **OK**.

Related concepts:

“Taxes for organizations” on page 371

You can specify default tax general ledger accounts and tax codes to calculate the amount of tax that is due on a purchase requisition, request for quotation, purchase order, or invoice. You can also specify the order in which tax codes are used in calculations.

Related reference:

“Customization options for applications” on page 369

There are many options and settings that you can customize for your applications in the Organizations application. They include work order options, inventory, purchasing, asset, and preventive maintenance (PM) options. You can access the options from the Select Action menu.

Specifying options for invoices

You can specify the invoice date as the start of the financial period. The invoice date and the general ledger posting date are matched, because the general ledger posting date is replaced with the invoice date. You can also specify that financial period validation is performed each time an invoice is saved.

About this task

If the value specified in the **Invoice Date** field is not a valid financial period, the value in the **G/L Posting Date** field is updated to the start date of the next open financial period.

For consignment invoices, you can specify the maximum number of lines that can be included on the invoice for the vendor that is replenishing the consignment items. When the default maximum number of lines specified is reached, a consignment invoice is automatically created.

Procedure

1. In the Organizations application, select the organization for which you want to set invoice options.
2. From the **Select Action** menu, select **Purchasing Options > Invoice Options**.
3. In the Invoice Options window, specify whether you want the **G/L Posting Date** field to be updated with the date in the **Invoice Date** field.
4. Specify whether you want the financial period validated each time an invoice is saved. Clear this option if you want the financial period validated when the invoice is approved.
5. For consignment types of invoices, specify the maximum number of lines that can be included on the invoice for a specific vendor. The default maximum number of lines is 40.
6. Click **OK**.

Related concepts:

“Taxes for organizations” on page 371

You can specify default tax general ledger accounts and tax codes to calculate the amount of tax that is due on a purchase requisition, request for quotation, purchase order, or invoice. You can also specify the order in which tax codes are used in calculations.

Related reference:

“Customization options for applications” on page 369

There are many options and settings that you can customize for your applications in the Organizations application. They include work order options, inventory, purchasing, asset, and preventive maintenance (PM) options. You can access the options from the Select Action menu.

Specifying autonumbering for applications

You can specify autonumber seeds and prefixes for applications at the system, organization, site, and set levels so that all sets in an organization use the same numbering sequence.

About this task

To create an autonumber or to attach an autonumber to a field, use the Database Configuration application.

Procedure

1. In the Organizations application, select the organization for which you want to set up autonumbering.
2. From the Select Action menu, select **System Level, Set Level, Organization Level, or Site Level**.
3. Specify the values for the appropriate level:
 - For the system level, set level, and organization level, specify or change values for the seed.
 - For the site level, specify or change values for the autonumber name and seed.
4. Optional: To have the autonumber preceded by a standard value, specify a value for prefix. For example, IN (for incident) can precede the autonumber.
5. Click **OK**.

Related concepts:

“Autonumbering” on page 367

When you set up autonumbering and a user creates a record, the record IDs or other specified fields increment by one.

“Application levels and data storage” on page 365

Application data is stored at different levels. The separation of data allows you to reuse as much data as possible across organizations and sites without breaching best practices in security and authorization.

Specifying autonumbering for special order items

Special order items are not kept in stock in the inventory, therefore, they do not have inventory item numbers. You must order these items by description.

Procedure

1. In the Organizations application, select the organization for which you want to specify autonumbering.

2. From the **Select Action** menu, select **Purchasing Options > PO Options**.
3. In the PO Options window, select **Allow the Generation of Special Order Items**.
4. Click **OK** and save your changes.

Related concepts:

“Autonumbering” on page 367

When you set up autonumbering and a user creates a record, the record IDs or other specified fields increment by one.

“Application levels and data storage” on page 365

Application data is stored at different levels. The separation of data allows you to reuse as much data as possible across organizations and sites without breaching best practices in security and authorization.

Displaying user messages

You can specify when users are prompted by user interface (UI) messages to provide them with information.

About this task

For assets that are down, if downtime was already reported, the message does not appear. Duplicate problem messages are displayed by default for assets and locations.

Procedure

1. In the Organizations application, select the relevant organization.
2. From the **Select Action** menu, select **Work Order Options > Other Organization Options**.
3. Specify the details you want to display to users:
 - To display asset status information after a user completes or closes a work order, and when the asset status is still down, select the **Display Downtime Report Prompt upon WO Completion for Asset in a 'Down' Status** check box.
 - To display a warranty message when a user enters an asset under warranty on a work order, select the **Display Warranty Status** check box.
 - To display a duplicate problem message when a duplicate problem is reported on an asset or on a location, select the **On Asset** or **On Location** check boxes.
4. Click **OK**.

Chapter 18. Managing calendars

You can create and change calendars for organizations and for sites. You can create calendars to define working times as well as shifts. A calendar can also specify non-working time, such as weekends, holidays, and shutdowns.

Calendars overview

Calendars are shared entities that define the framework for shifts, holidays, and so on, for organizations. A calendar can also specify non-working time, such as weekends, holidays, and shutdowns.

Shift patterns for calendars

A shift defines working time that is not specific to a date. You choose the working days, then you designate the start time and end time for work.

Once a shift is defined, you can apply it to a calendar. You can then apply the calendar to person, location, asset, and other records to specify working time.

Most shift patterns use a seven day pattern. Most patterns start with Sunday, or are multiples of seven, such as 14 days or 21 days, with a Monday start day. In some companies, there are unique circumstances for which a five day or other pattern might be used. You can create shift definitions that do not reflect the typical working time at your company, but would be useful for special work situations.

You use the **Define Pattern** action to create shift patterns. The number of days in the pattern specifies the block of days that repeat. If you use a number that is not a multiple of 7, the pattern does not repeat on the same days of the week. For example, with a 15-day pattern of 10 days on and 5 days off, the second instance of the shift starts on a different day than the first.

You can create a shift called first, which has the following properties:

- Working days are Monday through Friday
- Work starts at 7:00 a.m. and ends at 3:00 p.m.
- Work hours for the day total 8

Exceptions to the standard calendar

Information for individuals, such as vacation days, sick leave, personal time, and overtime, is not stored on the main calendar record.

You can use the following applications and icons to enter exceptions to the standard calendar:

People

Modify Person Availability

Assignment Manager

Modify Availability

The system combines the standard calendar assignments and the exceptions to determine the availability of a person for a given day, shift, and so on.

Working with calendars

You can create multiple calendars for an organization to manage shifts, working time, and non-working time.

Creating calendars

You can create calendars to define working time with a start date and an end date, as well as shifts. A calendar can also specify non-working time, such as weekends, holidays, and shutdowns.

Before you begin

Your default insert site must be in the organization for which you want to create the calendar. To check and change your default insert site, use the **Profile > Default Information** link in the toolbar.

About this task

The calendars that you create are for the organization of your default insert site.

Procedure

1. In the Calendars application, click **New Calendar**.
2. Type a name for the calendar.
3. Optional: Type a description for the calendar.
4. Specify the dates on which you want the calendar to start and end.
5. Save your changes.

What to do next

You created a calendar without shifts, work periods, holidays, or other details. Use the **Define/Apply Shifts** action to add shifts and the **Define/Apply Non-Working Time** action to add holidays and other non-working time. Use the People application to change person availability and the Assignment Manager application to change availability. In the calendar header, you can change the calendar description, start date, and end date.

Related concepts:

“Shift patterns for calendars” on page 379

A shift defines working time that is not specific to a date. You choose the working days, then you designate the start time and end time for work.

Specifying shifts in calendars

To customize shifts for your business needs, you can define shifts and apply them to calendars.

Procedure

1. In the Calendars application, select the calendar for which you want to define a new shift.
2. From the **Select Action** menu, select **Define/Apply Shifts**.
3. In the Define/Apply Shifts window, click **New Row**.
4. Type a name for the shift.
5. Optional: Type a description for the shift.
6. Specify the day on which you want the pattern to begin.

7. Type the number of days in the pattern. For example, the typical weekly patterns are seven days long, but you might need a 10-day, 14-day, or other pattern.
8. Optional: Define a pattern and apply the shifts to the calendar.
9. Click **OK**. The Define/Apply Shifts window opens with the details in the Shift Pattern table window.

Related concepts:

“Shift patterns for calendars” on page 379

A shift defines working time that is not specific to a date. You choose the working days, then you designate the start time and end time for work.

Applying shifts to calendars

In the Calendars application, you can specify the range of dates over which you want to apply shifts. You can apply the shifts to the entire period defined for the calendar or for a shorter period.

Procedure

1. In the Calendars application, select the calendar for which you want to apply a shift.
2. From the **Select Action** menu, select **Define/Apply Shifts**.
3. In the Define/Apply Shifts window, click **Apply Shift(s)**.
4. In the Apply Shifts with Range window, select one of the following ranges.

Option	Description
Entire Calendar	The start date and end dates reflect the dates specified for the calendar, but are read-only.
Selected Dates	The start date and end dates defined for the calendar are inserted. Change these dates, if necessary.

5. Click **OK**.

Related concepts:

“Shift patterns for calendars” on page 379

A shift defines working time that is not specific to a date. You choose the working days, then you designate the start time and end time for work.

Specifying shift patterns in calendars

To customize shift patterns for your business needs, you can specify the hours in a shift for each day in a pattern.

Before you begin

After you create a shift, you can specify a shift pattern.

Procedure

1. In the Calendars application, select the calendar for which you want to apply a shift.
2. From the **Select Action** menu, select **Define/Apply Shifts**.
3. In the Define/Apply Shifts window, click **Define Pattern**.

4. For the first work day in the pattern, specify the time the work period on that day starts and ends. When you move the cursor to the **Work Hours** field, the work hours are calculated as the interval between the two times.
5. Optional: Specify a different value for work hours.
6. Specify the start time and end time for the remaining work days. To copy the values from one day to the other days, click **Fill Out Work Days Data**.
7. Click **OK**.

Related concepts:

“Shift patterns for calendars” on page 379

A shift defines working time that is not specific to a date. You choose the working days, then you designate the start time and end time for work.

Copying calendars

The easiest way to create a calendar that is like an existing one is to copy the existing one and change the copy as needed.

Procedure

1. In the Calendars application, select the calendar that you want to copy.
2. From the **Select Action** menu, select **Duplicate Calendar**.
3. Type a name for the new calendar.
4. Optional: Type a description for the calendar. The default is the description on the calendar that you copied.
5. Save your changes. When you save the new calendar, all shift and non-working time information is copied from the old calendar.

Deleting calendars

If you are no longer using a calendar, you can delete it.

About this task

You cannot delete a calendar if it is used on any of the following types of records:

- Asset
- Asset status
- Personal calendar
- Service level agreement
- Preventive maintenance
- Job plan
- Location
- Work order

Procedure

1. In the Calendars application, select the calendar that you want to delete. You can select only one calendar at a time for deletion.
2. From the **Select Action** menu, select **Delete Calendar**.
3. Click **Yes**.

Establishing work periods

You can create work periods for a given date and then customize the work period to suit your business needs. You can also modify work periods, and specify non-working time for work periods.

Creating work periods

A work period represents a single shift applied to a specific date. You can create work periods for a given date and then customize the work period to suit your business needs.

Procedure

1. In the Calendars application, select the calendar for which you want to create work periods.
2. Click the **Work Periods** tab. All work periods for the entire calendar display, including all dates with defined work hours or defined non-working time.
3. Click **New Row**.
4. Specify values for work date, shift, start time, end time, and work hours.
5. Optional: Add a note.
6. Save your changes.

Changing work periods

A work period represents a single shift applied to a specific date. To customize work periods for your business needs, you can change work periods for a given date.

About this task

You cannot change a non-working time period. However, you can delete it.

You also can change work periods from the **Calendar** tab. On the date for which you want to change working time, click the hours value and make the necessary changes.

Procedure

1. In the Calendars application, select the calendar that you want to change.
2. Click the **Work Periods** tab. All work periods for the entire calendar are displayed, including all dates with defined work hours or defined non-working time.
3. Find the row with the work period that you want to change.
4. Change the start time, end time, work hours, and notes.
5. Save your changes.

Example

If you have a corporate calendar, and you want to change the work period for a single shift for a single date because the company is closing early, you change the shift for that date.

Specifying non-working time for work periods

You can define non-working time, such as holidays and shut downs, or any other non-working time that you want to apply to a work period.

Procedure

1. In the Calendars application, select the calendar for which you want to define non-working time.
2. From the **Select Action** menu, select **Define/Apply Non-Working Time**.
3. In the Define/Apply Non-Working Time window, click **New Row**.
4. Optional: Type a description for the non-working time.
5. Specify a start date and end date. For non-working time that occurs on one day, the start date and the end date are the same.
6. Specify a value for type.
7. Click **Apply**.
8. Click **OK**.

Related concepts:

"Exceptions to the standard calendar" on page 379

Information for individuals, such as vacation days, sick leave, personal time, and overtime, is not stored on the main calendar record.

Chapter 19. Managing classifications

You can use classifications to simplify the task of managing and of retrieving historical data from other applications. You can use classifications to align with external standards, such as vendor standards and industry standards. You can also use classifications to define the escalation path for incidents.

Classifications overview

Classifications identify and characterize similar objects. A building, a notebook computer, and a centrifugal pump are types of classifications. A classification can also describe an event, such as a broken window or a hard disk failure. You use the Classifications application to create classifications and to establish classification hierarchies.

Related tasks:

“Creating classifications” on page 390

You classify information to categorize it logically, so that the information is easier to find. You can classify different types of records, such as location records, asset records, item records, and work order records.

“Modifying classifications” on page 392

You can change classifications by modifying the information for the classification. You can also modify a classification by adding a classification to an existing classification to create a hierarchy for the top-level classification.

Classification paths and hierarchies

You can develop a classification structure by joining two or more classifications into a hierarchy. In a classification hierarchy, the next level up is called the parent, and the next level down is called the child. Each child (which is also a classification) can become a parent and can have its own child levels. This pattern can continue indefinitely. The structure from the top-level parent to the child is called the classification path.

You can also use a numbering system in which a number describes the position of the classification in the classification structure. The number of digits describes the level in the hierarchy. Numbering systems are suitable only if they are not subject to change, because the identifier of the classification (PUMP, 30612456, ROTARY) cannot be changed. The identifier can be deleted only if it is not used in a classification hierarchy.

Example of a request for information

Someone calls the service desk to get information about the facilities. The service desk uses classifications to categorize the ticket as an informational request. The ticket is associated with the classification path INFORMATIONAL QUESTIONS \ FACILITIES \ WHERE DO I.

10 minutes later, the service desk receives a call about health coverage. The service desk can use the same classification structure to categorize both incidents. The following figure shows a classification structure that uses the example of a request for information.




Figure 10. Classification structure - example of a request for information

Example of a request for software installation

Someone contacts the service desk and requests a Microsoft Windows XP installation. A service request and change ticket are created. The ticket is associated with the classification path IT \ SOFTWARE \ INSTALLS \ WINDOWS XP Professional.

10 minutes later, the service desk receives a printer-related call. The service desk can use the same classification structure to classify the printer incident ticket as was used to categorize the earlier request.

The following figure shows a classification structure that uses the example of a request for software installation.




Figure 11. Classification structure - example of a request for software installation

Generate Description option

On the **Classifications** tab of the Classifications application, the **Generate Description** option and the **Use Classification** option provide varying results. The results depend on the option that you select.

The following table describes the results of using these options.

Option	Result
Select Generate Description only	The description that is provided in the Description field for the record, such as an asset record, that uses the classification is overwritten. The description is replaced with the attribute values that are entered in the Specifications tab of the record. Included in the description are the units of measure that are specified in the Attributes window of the Classifications tab. You can view generated descriptions in the application with which the record is associated, such as the Assets application.
Select both Generate Description and Use Classification	The description in the application with which the record is associated is overwritten with the identifier for the classification. The description is also overwritten with the values that are entered in the Specifications tab, and the units of measure that are specified in the Attributes window of the Classifications tab.

Associations of records with classifications

You can associate types of records, such as assets, items, or solutions, with specific classifications. You modify the Use With settings in the Classifications application to associate records with classifications.

The Select Value windows that display the classification can also be associated by the Use With settings.

You can tailor the attributes of a classification for each object that you want to use with it. You use the Attributes window to associate the attributes of classifications with specific objects.

Classification searches

You can search classification structures and attributes that have associated values. You can search for values in records that you created for items, locations, configuration items, work orders, tickets, and so on.

Example of a classification search

Your company created classifications for the following fields:

- Name
- Affected by (name)
- Asset
- Location
- Service Groups
- Service

When you create a service request record, you enter values in the **Name** field, and in the **Asset** field. You also enter values in the **Location** field, the **Service Groups**

field, and the **Service** field. You can use the navigation tree in the Classification Search window to view the values that are associated with the classification structures and their attributes.

The classification hierarchy that is shown varies depending on the application or the field from which you access the window. The navigation tree shows classifications that can be used with a specific type of record.

Actual and authorized configuration item classifications

You can define associations between actual configuration item (CI) classifications and authorized CI classifications.

Actual configuration items

An actual configuration item is an item that is loaded into the Actual Configuration Items application from the discovery engine using the IBM Tivoli Integration Composer. You can load configuration items into the database from a discovery engine or an asset management database. First load your configuration item data into the discovery engine using the bulk loader interface.

Authorized configuration items

An authorized configuration item is an item that conforms to rules and relationships that you define in the Configuration Items application. These configuration items are items that you want to maintain under configuration management and change control.

When you create an authorized configuration item, it can be linked to an actual configuration item record. You can also create an authorized configuration item that is independent of any actual configuration item record, because not all configuration items need to work with the discovery engine. You can create a link between the authorized configuration item and the actual configuration item in the Configuration Items and Actual Configuration Items applications.

Actual configuration item classifications and authorized configuration item classifications

You use the Classifications application to define associations between actual configuration item (CI) classifications and authorized CI classifications. You can create an authorized CI from an actual CI record when you define these associations. You can include the number of attributes that you need for configuration management and change control.

You create an authorized classification hierarchy that maps to an actual classification hierarchy. You typically create partial mappings between the two classifications, rather than a complete mirror image. When creating an authorized CI from an actual CI you can create authorized configuration items. These items are created for the child classifications that have a containment relationship with the actual CI.

You use the discovery tool to discover actual configuration items in your environment. The number of actual configuration items that are discovered by the discovery tool is likely to be greater than the number of authorized configuration items for which you want to create records or to use. By mapping one hierarchy to another, you can limit the number of actual configuration items that you link to authorized configuration items. For example, an actual configuration item hierarchy has five levels. You can create authorized configuration item records for the top three levels of actual configuration items.

Attributes of classifications

You use attributes to define the characteristics of classifications. For example, for a classification of pump, you can have an attribute of horse power. For a classification path of PROBLEM\COMPUTER, you can have an attribute of memory size.

You can further define attributes by domain. A classification attribute can be tied to an ALN domain, to a numeric domain, or to a table domain. You can validate an attribute against the values in a particular database column by using a table domain.

You define domains using the Domains application.

Groupings of attributes

You can group attributes into sections so that you can use the same attribute multiple times.

Example of an attribute grouping

For example, you define a pipe in the system as an asset. The pipe has the following characteristics:

- It is 80 feet (~25 m) long.
- It contains 10 sections of equivalent length.
- It has an interior diameter that narrows at one end of the pipe.

Because of the taper, the walls of the pipe must be thicker at the narrow end to withstand the higher pressure. Each section has a different average interior diameter and wall thickness. Therefore, the attribute is the interior diameter.

Apply Down Hierarchy option

When you add an attribute to a classification, you can use the **Apply Down Hierarchy** option. When you check this box, the attribute that is assigned to that classification is added to the classifications that are below it in the hierarchy.

The **Inherited From** field indicates that the attribute is associated with a classification that is higher up in the hierarchy. The **Inherited From** field displays the classification where an attribute originated.

Classifications planning

You create classifications to retrieve information later. Before you create a classification, determine the information that you want to retrieve. You can base your classification structure on how you group objects in your enterprise. You can also use classifications to define the escalation path for incidents.

Types of records to be classified

You can use the Classifications application to store information about different types of records. You can classify application records and search for the records that you classified.

You can classify the following types of records:

- Assets
- Items
- Locations

- Solutions
- Tickets (service requests, incidents, and problems)
- Work orders (activities, changes, and releases)

Categories of items for reporting

You can organize your classification structures into categories or groups of items for reporting. You can create unlimited classification levels.

You can organize information into top-level categories, such as the following categories:

- Information Technology assets
- Production assets
- Facility assets
- Fleet assets

You can work from the top-level categories into the more detailed levels. For example, fleet assets can contain 18-wheel trucks, and sales fleet cars.

Another example is to categorize the service desk call for a software installation differently than the service desk call for information. This type of categorization helps to determine how many customers complained about a software installation as opposed to how many calls were about health coverage.

Industry codes

You can apply industry codes when you create classifications. You can get a list of the codes that you frequently use in your enterprise, and apply them to your classifications.

Example of using an industry code

If your company does vehicle or code maintenance, you can use the vehicle maintenance industry standard codes. In the United States, most mechanics know that an oil change is code 42-3-2. You can use the same code when you create a classification for an oil change, or use the code that applies in your country.

Working with classifications

You can create classifications, and specify details about the attributes of classifications. You can search classification structures and attributes that have associated values when you use any record type that can be classified. You can also define associations between actual configuration item (CI) classifications and authorized CI classifications.

Creating classifications

You classify information to categorize it logically, so that the information is easier to find. You can classify different types of records, such as location records, asset records, item records, and work order records.

About this task

To make a classification available for every organization and site, leave the Site and Organization fields blank.

Procedure

1. From the application toolbar of the Classifications application, click **New Classification**.
2. In the **Classification** field, specify a classification.
3. Specify the information for the new classification.
4. Optional: To specify the records that you want to use with the classification, click **New Row** in the Use With table and add one or more records.
5. To add child classifications, click **New Row** in the Children table.
6. Optional: To add attributes, insert rows and complete the fields in the Attributes table. For most of the fields, click **Select Value** to select from existing values. Alternatively, you can create values, as with the **Classification** field.
7. Click **Save Classification**.

Related concepts:

“Classifications overview” on page 385

Classifications identify and characterize similar objects. A building, a notebook computer, and a centrifugal pump are types of classifications. A classification can also describe an event, such as a broken window or a hard disk failure. You use the Classifications application to create classifications and to establish classification hierarchies.

“Generate Description option” on page 386

On the **Classifications** tab of the Classifications application, the **Generate Description** option and the **Use Classification** option provide varying results. The results depend on the option that you select.

“Classifications planning” on page 389

You create classifications to retrieve information later. Before you create a classification, determine the information that you want to retrieve. You can base your classification structure on how you group objects in your enterprise. You can also use classifications to define the escalation path for incidents.

Associating attributes with records

You can specify detailed information about the attributes of a classification. You can enter attribute details about each record that you want to use with a classification.

Procedure

1. From the Classifications application, click the **Classifications** tab.
2. In the Attributes section, click the **Use with Object Detail** icon on the right side of the row for the attribute whose details you want to enter.
3. In the Use With Object Detail window, specify details for the attribute.

Field or column name	Description
Use With Object	The record for which you enabled the classification.

Field or column name	Description
Sequence	The order in which you want the attribute to appear in the list of attributes on the Specifications tab of the related application. For example, the attribute capacity is one of three attributes of a classification that you enabled for the ASSET and WORKORDER records. In the Assets application, the attribute capacity can appear first in the list of attributes. In the Work Order Tracking application, the attribute capacity can appear third.
Mandatory	Makes the attribute a required field in the related application. For example, you can make it a required field in the Assets application, but not in the Work Order Tracking application.
Used in Description Generation	The description of the attribute appears in the description of the classification in the related application.
Used in Specifications	The attribute appears on the Specifications tab of the related application.

4. Optional: Specify a default value for the attribute:
 - a. In the row for the attribute, click **View Details**, to the left of the Use With Object name column.
 - b. In the Details area, specify a default value based on the data type of the attribute.
5. Click **OK**.

Related concepts:

“Attributes of classifications” on page 389

You use attributes to define the characteristics of classifications. For example, for a classification of pump, you can have an attribute of horse power. For a classification path of PROBLEM\COMPUTER, you can have an attribute of memory size.

“Groupings of attributes” on page 389

You can group attributes into sections so that you can use the same attribute multiple times.

“Classifications planning” on page 389

You create classifications to retrieve information later. Before you create a classification, determine the information that you want to retrieve. You can base your classification structure on how you group objects in your enterprise. You can also use classifications to define the escalation path for incidents.

Modifying classifications

You can change classifications by modifying the information for the classification. You can also modify a classification by adding a classification to an existing classification to create a hierarchy for the top-level classification.

Procedure

1. From the Classifications application, click the **Classifications** tab.
2. From the **Select Action** menu, select **Add/Modify Properties > Classifications**.
3. To add a classification to an existing classification, perform the following steps:

- a. Click **New Row**.
 - b. Type a unique classification name.
 - c. Type a description, an organization, and a site for the classification.
 - d. Repeat this process for each classification that you want to add.
4. To modify a classification:
 - a. Select the classification that you want to modify.
 - b. Modify the information for the classification.
5. Click **OK**.

Related concepts:

“Classifications overview” on page 385

Classifications identify and characterize similar objects. A building, a notebook computer, and a centrifugal pump are types of classifications. A classification can also describe an event, such as a broken window or a hard disk failure. You use the Classifications application to create classifications and to establish classification hierarchies.

“Generate Description option” on page 386

On the **Classifications** tab of the Classifications application, the **Generate Description** option and the **Use Classification** option provide varying results. The results depend on the option that you select.

Modifying attributes

You use attributes to define the characteristics of classifications. You can modify the properties of attributes. You can then associate the attributes with classifications.

Procedure

1. From the Classifications application, click the **Classifications** tab.
2. From the **Select Action** menu of the **Classifications** tab, select **Add/Modify Properties > Attributes**.
3. From the Attributes area of the **Classifications** tab, find the attribute that you want to modify.
4. Edit the appropriate fields.
5. Repeat this process for each attribute that you want to modify, and click **OK**.

Adding attributes

You can add attributes. You can then associate the attributes with classifications. You can use attributes with many record types, such as asset records, location records, and item records.

Procedure

1. From the Classifications application, click the **Classifications** tab.
2. From the **Select Action** menu, select **Add/Modify Properties > Attributes**.
3. Add a row in the Attributes area, and specify the information for the new attribute.
4. Repeat this process for each attribute that you want to add, and click **OK**.

Related concepts:

“Attributes of classifications” on page 389

You use attributes to define the characteristics of classifications. For example, for a classification of pump, you can have an attribute of horse power. For a classification path of PROBLEM\COMPUTER, you can have an attribute of memory size.

“Groupings of attributes” on page 389

You can group attributes into sections so that you can use the same attribute multiple times.

“Apply Down Hierarchy option” on page 389

When you add an attribute to a classification, you can use the **Apply Down Hierarchy** option. When you check this box, the attribute that is assigned to that classification is added to the classifications that are below it in the hierarchy.

Searching for classifications from application records

You can search classification structures and attributes with associated values when you use any record type that can be classified. Record types include assets, items, locations, configuration items, work orders, and tickets.

Procedure

1. From the application associated with the type of record, click the **Detail Menu** button and choose **Classification**.
2. Select a classification category in the Classification Search tree on the left, and drill down to find the record.
3. Select a record identifier in the window below the Classification Search tree to return a value to the field.

Related concepts:

“Classification searches” on page 387

You can search classification structures and attributes that have associated values. You can search for values in records that you created for items, locations, configuration items, work orders, tickets, and so on.

Defining associations between actual and authorized configuration item classifications

You can define associations between actual configuration item (CI) classifications and authorized CI classifications. When you map associations, you create an authorized CI record from an actual CI record.

Before you begin

The classification must be at the top level of the hierarchy. You define the top level for a classification when you create the classification.

Additionally, in the Use With window, the applications must include configuration items.

About this task

For each actual CI classification hierarchy that you use to import actual configuration items, you can create one or more authorized classification hierarchies.

For each of the authorized CI hierarchies, you use the Manage CI Hierarchies window to map the links between authorized and actual classifications. Mapping includes the parent-child classification levels. At each level of the classification, you can also specify the relationships between source classifications and target classifications.

Procedure

1. From the **List** tab of the Classifications application, select the CI classification hierarchy that you want to manage.
2. From the **Select Action** menu, select **Manage CI Hierarchies**.
3. In the **Actual CI Classification** field at the top of the Manage CI Hierarchies window, click **Detail Menu** and select **Classify** to select an actual classification.
4. In the Children window, select a row.
5. In the **Actual CI Classification** field, specify the actual CI classification that you want to associate with the authorized child classification.
6. Click **Detail Menu** and select **Classify** to select a classification.
7. In the Relationships window, specify the relationships that you want to define for the selected child classification. To add a relationship, perform the following steps:
 - a. Click **New Row**.
 - b. Specify the source and target classifications, the relationship, and other information.
 - c. Click **Select Relation Rules** to select from a list of relationship rules.
8. Optional: Repeat steps 4, 5, and 6 to map actual classifications to other child authorized classifications.
9. Click **OK**.

Results

After you map the classifications, you use them in the Actual Configuration Items application to create configuration items from actual configuration items. When you define mappings between authorized and actual configuration items, you can create many configuration item records at the same time. You link entire or partial hierarchies instead of individual configuration items. You create records for the child classifications and for the top-level classification.

Related concepts:

“Actual configuration item classifications and authorized configuration item classifications” on page 388

You use the Classifications application to define associations between actual configuration item (CI) classifications and authorized CI classifications. You can create an authorized CI from an actual CI record when you define these associations. You can include the number of attributes that you need for configuration management and change control.

Adding units of measure

You can specify or change the units of measure that you use when ordering and issuing items. Any units of measure that you add become available in the product for use with all items.

About this task

If the units of measure do not match when you receive or transfer an item into a storeroom, the product looks for a conversion ratio to determine the stocking balance. You associate the order and issue units to an item when you add the item to a storeroom. You can access the Add/Modify Units of Measure window from more than one application.

Procedure

1. From your application, display a record.
2. From the **Select Action** menu, select **Unit of Measure and Conversion > Add/Modify Units of Measure**.
3. Click **New Row**.
4. Specify a unit of measure.
5. Type an abbreviation for this unit of measure. If configured to do so, the product uses the abbreviations along with the classification of an asset to create descriptions for assets, locations, and items.
6. Type a description for this unit of measure.
7. Click **New Row** to specify additional units of measure.
8. Click **OK**.

Configuring the signature option to enable classification images

Signature options, which are defined in the Security Groups application, are used control access to functions in applications. By using signature options, you can grant access to specific groups of users. Under the Classifications application, you can select images that represent asset classifications that are used, for instance, in the Assets application topology view.

Procedure

1. In the **List** tab of the Security Groups application, you select the group to be granted access to the relationship functions. For example, MAXEVERYONE for all Maximo users.
2. Under the **Applications** tab, in the **Description** column of the top panel, click **Classifications**.
3. Scroll down to the third panel, listing Classifications application functions, select the **Grant Access** check box on the **Add/Modify Image** line, and **Save**.

Adding images to the asset topology view

In the **Topology** tab of the Assets application, assets can be represented by images based on their classification. These images are assigned to assets in the Classification application.

Procedure

1. In the Classifications application, select a classification.
2. Select the **Add/Modify Image** action. The image must be a GIF or JPG file.
3. Enter the file location or click **Browse** to navigate to it, click **OK** and then **Save**.

Chapter 20. Managing charts of accounts

You use the Chart of Accounts application to set up default general ledger accounts and resource codes for standard accounting functions.

Chart of accounts overview

You can set up default general ledger accounts and resource codes for standard accounting functions. You create accounts and resource codes to correspond with accounts that you use in your external accounting system. You can also set up financial periods, and specify general validation options.

General ledger account codes

You define values for the different components, and link together the component values to create general ledger account codes. You use the general ledger accounts for specific financial tracking purposes. You can also specify dates for general ledger accounts to be active or to expire.

A general ledger account code consists of several components (or segments), that are separated by delimiters, such as 6000-200-350. A component that has not been assigned a value is represented by placeholder characters, such as 6000-???-350.

You define the format of the account code in the following ways:

- You can define the number, length, and data type of components, whether the components are required, and any delimiter, in the Database Configuration application. The account structures defined in Database Configuration are the default, system-level structures.
- You can define the format of the account code to be specific to an organization. You can define the number, length, and data type of components, whether the components are required, and any delimiter, using the Chart of Accounts application.
- You can specify a default placeholder character in the Database Configuration application. You can specify an organization-specific placeholder character in the Organizations application.

You can specify the validation rules for the following component combinations for the general ledger account codes that users can use:

- A combination of existing components
- Only component combinations that are specified as general ledger accounts

You also can download account codes from your accounting system. A generic financial application programming interface and several product-specific application programming interfaces are provided. These application programming interfaces let the system interface with financial software, such as Oracle® and SAP®.

Related tasks:

“Creating general ledger account codes” on page 399

You can use general ledger account codes to enable financial reporting. You specify the component values for general ledger accounts. You can also specify a date range for an account to be active. For example, you can specify that a general ledger account for a specific craft is active from 02/25/2010 through 02/25/2011, and expires on 02/26/2011.

“Changing general ledger account codes” on page 400

After you create a general ledger account, you can change specific information. You might want to activate an account code or set the expiration date on an account code.

Organizational default accounts for general ledgers

In the Chart of Accounts application, you use organizational default accounts as holding accounts for financial considerations.

There are three organizational default accounts:

- Global rotating suspense account - You use rotating suspense accounts to hold the accumulated cost of repairs for rotating equipment.
- Global ticket account - You use the global ticket account when a ticket for a service request is created and no other account is available.
- Tool control account - You use the tool control account when a transaction involves a tool and no other account is available.

In the Chart of Accounts application, you use the **Organization Default Accounts** action to define default accounts.

Merge of general ledger accounts

When a general ledger account field is not uniquely specified, the general ledger components are merged. Generating work orders and other kinds of transactions requires choosing among component values. The product invokes a set of rules on how to handle them.

General ledger accounts are merged component by component. Defined components always supersede an undefined component. For example, if the first component of one account code is 6000 and the first component of the other account code is a placeholder (????), the merged first component is 6000.

When general ledger account codes merge, the result can be an unspecified account code. In this case, to complete the merge, complete one of the following actions:

- To use the merged account code, an authorized user must establish the account code in the Chart of Accounts application. The account code must be established before you can proceed with the transaction.
- If you do not use the merged account code, then ensure that the transaction is valid.

Resource codes for general ledgers

A resource code typically consists of one component of the account code.

There are three kinds of resource codes:

- Labor resource codes - General ledger account code components that track whether labor used in a transaction is internal labor or external labor.
- Tool resource codes - General ledger account code components that track whether a tool used in a transaction belongs to the enterprise (internal) or an external vendor.
- Inventory resource codes - General ledger account code components to attach to inventory items that are used in transactions.

Related tasks:

“Specifying resource codes for general ledgers” on page 406

You can specify default values for different types of labor resources, tool resources, and inventory items. These codes are then used to determine accounting charges.

Inactive component values

In the Chart of Accounts application, you can deactivate component values. Inactive components are not available to general ledger accounts that you add to records. General ledger accounts that use an inactive component become inactive also. You can also set component values to expire by a specific date, which renders them inactive.

When you deactivate a general ledger component value, no change is made to the general ledger accounts on existing records that use that value.

Example

A work order uses a cost center component with a value of 6250. You deactivate this cost center. The value is not available to use in new general ledger accounts. All general ledger accounts that use the cost center also become inactive. The existing work order still uses cost center 6250.

Working with chart of accounts

You can set up default general ledger accounts and resource codes for standard accounting functions.

Working with general ledger accounts

You can change the configuration of general ledger account structures to be specific to your organization. You can create component codes and resource codes to correspond with accounts that you use in your external accounting system. You can also set up financial periods, and specify general validation options.

Creating general ledger account codes

You can use general ledger account codes to enable financial reporting. You specify the component values for general ledger accounts. You can also specify a date range for an account to be active. For example, you can specify that a general ledger account for a specific craft is active from 02/25/2010 through 02/25/2011, and expires on 02/26/2011.

About this task

If the date set in the **Expiration Date** field is in the future, the general ledger account remains active until the specified date is reached. If no expiration date is specified, the general ledger account does not expire.

Procedure

1. In the Chart of Accounts application, select the organization for which you want to create a general ledger account code.
2. Click **New Row**.
3. Specify the component values for the general ledger account value.
4. Click **OK**.
5. Optional: Change the long description. If you change the descriptions for component values, the general ledger account description is not updated. Update the general ledger account description manually.
6. Optional: Specify the appropriate type code.
7. The default for the **Active Date** field is the date that the general ledger account is created. To activate the account at a different date, specify a future date. The general ledger account remains inactive until the date specified is reached. The read-only **Active** check box shows as cleared to indicate that the account is not active.
8. If you want the general ledger account to expire by a specific date, specify a date in the **Expiration Date** field.
9. Save your changes.

Results

After you save your changes, you cannot edit the **GL Account** field.

Related concepts:

“General ledger account codes” on page 397

You define values for the different components, and link together the component values to create general ledger account codes. You use the general ledger accounts for specific financial tracking purposes. You can also specify dates for general ledger accounts to be active or to expire.

Changing general ledger account codes:

After you create a general ledger account, you can change specific information. You might want to activate an account code or set the expiration date on an account code.

About this task

If the date set in the **Expiration Date** field is in the future, the general ledger account remains active until that date is reached. If no expiration date is specified, the general ledger account never expires.

Procedure

1. In the Chart of Accounts application, select the organization for which you want to change the general ledger account code.
2. Click **View Details** for the code that you want to change.
3. Optional: Change the information in the **GL Account Description** field, the **Type** field, the **Active Date** field, and the **Expiration Date** field.
4. Save your changes.

Related concepts:

“General ledger account codes” on page 397

You define values for the different components, and link together the component values to create general ledger account codes. You use the general ledger accounts for specific financial tracking purposes. You can also specify dates for general ledger accounts to be active or to expire.

Deleting general ledger account codes:

As your business needs change, you can delete a general ledger account code at any time.

About this task

When you delete a general ledger account code, you cannot use the account code on new records. However, existing database records that used the deleted account code are not changed. Therefore, a deleted general ledger account code can still exist in previously created work orders, purchase orders, and so on.

Deleting general account codes can cause problems with the general accounting system that you integrate with the system.

Procedure

1. In the Chart of Accounts application, select the organization for which you want to delete a general ledger account code.
2. Click **Mark Row for Delete**.
3. Save your changes.

Creating general ledger component values

You define valid component values for general ledger account codes. You use these values when you select a general ledger account code.

Procedure

1. In the Chart of Accounts application, select the organization for which you want to create a general ledger component value.
2. From the **Select Action** menu, select **GL Component Maintenance**.
3. In the GL Component Maintenance window, select the component for which you want to create values.
4. Click **New Row**.
5. In the **GL Component Value** field, specify the value for the component. The value that you specify must be in the format that was specified in the Database Configuration application. The value in Database Configuration is the default, system-level format.
6. Type a description of the component value.
7. Optional: To deactivate the component value, clear the **Active** check box.
8. Click **OK**.

Changing component values in general ledger accounts

To adapt to business changes, you can change the description of a general ledger component value. You can also activate and deactivate components that are used in general ledger accounts.

About this task

If you change the status of a component value to inactive, all general ledger accounts that use that component inherit that status. The **Expiration Date** field for the associated general ledger account on the main Chart of Accounts window is set to the current date and time. The **Expiration Date** field is set to the current date and time only if the field was previously set to a future date, or if the field was left blank. If the field was set to a date in the past, that value is restored.

When you reactivate an account component, you reactivate the general ledger accounts that use the component. Additionally, the **Active Date** field on the main Chart of Accounts window is set to the current date and time if the **Active Date** previously specified was in the future. If the **Active Date** field was set to a date in the past, that value is restored.

The **Expiration Date** field displays the expiration date that was previously specified if the date set is in the future. If the expiration date previously specified is in the past, that value is restored. If an expiration date was not specified, the **Expiration Date** field is left blank.

Procedure

1. In the Chart of Accounts application, select the organization for which you want to change a component value.
2. From the **Select Action** menu, select **GL Component Maintenance**.
3. In the GL Component Maintenance window, click **View Details** for the component value that you want to change.
4. To change the component value, perform one of the following actions:
 - Change the description.
 - Change the active status of the component value in the **Active** check box. If you activate a component value that was previously deactivated, the Activate GL Accounts window displays. Click the **Activate Accounts** check box for the component that you are reactivating, then click **OK**.
 - Change additional information.
5. Click **OK**.

Deleting general ledger component values

To customize general ledger information for your business needs, you can delete a component value that is used in general ledger account codes.

Procedure

1. In the Chart of Accounts application, select the organization for which you want to delete a general ledger component value.
2. From the **Select Action** menu, select **GL Component Maintenance**.
3. In the GL Component Maintenance window, click **Mark Row for Delete** for the component value that you want to delete.
4. Click **OK**.

Example

After you delete a general ledger component value, general ledger account codes in existing database records that use that value are not changed. For example, a work order uses a cost center value of 6500. If you delete 6500, the work order keeps this number as the cost center.

Changing general ledger account structures

You can change the default component values for general ledger accounts to be specific to organizations within your enterprise. You might want organizations to have distinct general ledger account structures. For example, the lengths of the components might vary, and the components might be of different data types, depending on how the components are being used.

About this task

The first time the Add/Modify Account Structure window displays, the default configuration defined for the general ledger account displays. You use the Database Configuration application to define the default configuration at the system level. The changes made to organization-specific account structures do not affect the configuration specified at the system level.

You can define one account structure specific to each organization. When the account structure is defined, you can modify the attributes. The modified account structure does not affect existing data.

Procedure

1. In the Chart of Accounts application, select the organization for which you want to modify the general ledger account structure.
2. From the **Select Action** menu, select **Add/Modify Account Structure**.
3. In the Add/Modify Account Structure window, select the component for which you want to modify values.
4. In the **Component** field, specify the name of the component.
5. In the **Length** field, specify the length of the field. The length of the field cannot exceed the length defined at the system level.
6. In the **Type** field, specify the data type for the component. If the type defined at the system level is numeric, the type set at the organization level cannot be alphanumeric. If the type set at the system level is alphanumeric, the type set at the organization level can be numeric.
7. Optional: Click the **Required** check box if you want the component to be mandatory. If the value set at the system level is 1, the **Required** check box is read-only. If the value set at the system level is 0, this check box is modifiable for the organization.
8. Optional: Modify the **Screen Delimiter** to specify a different separator character between components.
9. Click **OK**.

Example

For example, you can create account structures for different organizations in separate geographic regions, as follows:

- Organization A is in Europe. Organization A has the account structure **xxx-xx-xxx** with component one and component two being required. Component two is of the alphanumeric type.
- Organization B is in North America. Organization B has the account structure **xxx-xxx-xx** with components one, two, and three being required. Component two is of the integer data type.

Updating databases for general ledger accounts

After you change a default general ledger account or resource code, you can update the database for one organization at a time.

Before you begin

When you update the database, ensure that no one is using the system. Historical records are not updated.

Procedure

1. In the Chart of Accounts application, select the organization whose database you want to update.
2. From the **Select Action** menu, select **Update Database**.
3. In the Update Database window, select one of the following updates:

Option	Description
Overwrite Blank Accounts Only	Use this option to overwrite blank GL Account fields
Overwrite Accounts With Old Defaults	Use this option to overwrite accounts that have not been updated since being inserted.
Overwrite All Accounts	Use this option to overwrite all relevant GL Account fields, including blank fields, with updated data.

4. Click **OK**.

Setting up accounts

You can set up organizational default accounts as holding accounts for financial considerations. You can also specify default general ledger accounts to be associated with company-related accounts, as well as external labor control accounts to be associated with general ledger accounts.

Setting up organization default accounts

You use organizational default accounts as holding accounts for financial considerations.

About this task

You can set up the following global general ledger accounts:

Global Rotating Suspense Account

You use rotating suspense accounts to hold the accumulated cost of repairs for rotating equipment.

Global Ticket Account

You use the default account when a ticket for a service request is created and no other account is available.

Tool Control Account

You use the default account when a transaction involves a tool and no other account is available.

Procedure

1. Open the Chart of Accounts application.

2. In the Organizations table window, select the organization for which you want to define the accounts.
3. From the **Select Action** menu, select **Organization Default Accounts**.
4. In the **Global Rotating Suspense Account** field, click **Select Value** and choose an account.
5. In the **Global Ticket Account** field, click **Select Value** and choose an account.
6. In the **Tool Control Account** field, click **Select Value** and choose an account.
7. Click **OK**.

Specifying company-related accounts for general ledgers

For types of companies that use payment, you can specify default general ledger accounts for company-related accounts.

About this task

The accounts include accounts received but not invoiced, account payable suspense, and accounts payable control account. The accounts are based on company types that are specified in the Companies application.

Procedure

1. In the Chart of Accounts application, select the organization for which you want to specify company-related accounts.
2. From the **Select Action** menu, select **Company Related Accounts**.
3. In the Company Related Accounts window, click **New Row**. You can add a new row only if there is an enterprise type that has not yet been associated with general ledger accounts.
4. Specify a type, received but not invoiced account information, accounts payable suspense account information, and accounts payable control account information.
5. Click **OK**.

Specifying external labor control accounts for general ledgers

You use external labor control accounts to set up default account codes for work performed by outside vendors.

Procedure

1. In the Chart of Accounts applications, select the organization for which you want to specify an external labor control account.
2. From the **Select Action** menu, select **External Labor Control Accounts**.
3. In the External Labor Control Accounts window, click **New Row**. You can add a new row only if there is a vendor that has not yet been associated with a general ledger account.
4. Specify values for the vendor and control account.
5. Click **OK**.

Specifying financial periods for general ledgers

You can segment accounting periods by specifying the start date and close date of a financial period.

Before you begin

You must define at least one financial period.

About this task

The transactions must occur during an open, valid financial period. The requirements of the accounting system that you use determines the format of the financial period.

Procedure

1. In the Chart of Accounts application, select the organization for which you want to define a financial period.
2. From the **Select Action** menu, select **Financial Periods**.
3. In the Financial Periods window, click **New Row**.
4. Enter the name or number for the period.
5. Specify start and end dates for the financial period: If there is no financial period, the current date and time displays in the **From** field. If financial periods exist, the **From** field displays the date and time shown in the **To** field of the most recent period.
6. Optional: Specify a close date. After this date, you cannot charge additional transactions to the accounting period.

Results

To prevent gaps and overlaps in time between contiguous periods, if you change the date of an existing period, the surrounding dates are reset. If you use financial periods, then a financial period stamp is added to all transactions when they are generated.

What to do next

You use the **Actual Close Date** field to close the financial period.

Closing financial periods

You can close a specific financial period to ensure that no more transactions are entered in that period.

Procedure

1. In the Chart of Accounts application, select the organization for which you want to close a financial period.
2. From the **Select Action** menu, select **Financial Periods**.
3. In the Financial Periods window, specify the actual close date for the financial period that you want to close. Your name is displayed in the **Closed By** field.

Results

After you close a financial period, financial transactions for that period are no longer accepted.

Specifying resource codes for general ledgers

You can specify default values for different types of labor resources, tool resources, and inventory items. These codes are then used to determine accounting charges.

Procedure

1. In the Chart of Accounts application, select the organization for which you want to define resource codes.

2. From the **Select Action** menu, select **Resource Codes**.
3. In the Resource Codes window, complete the information specific to the type of resource code that you are defining:

Option	Description
Labor resource codes	In the Internal field and External field, specify the labor resource code components.
Tool resource codes	In the Internal field and External field, specify the tool resource code components.
Inventory resource codes	1. Click New Row. 2. Specify values for the commodity group. 3. Optional: Specify a value for the inventory resource code component.

4. Click **OK**.

Related concepts:

“Resource codes for general ledgers” on page 398

A resource code typically consists of one component of the account code.

Specifying validation options

You use validation options to specify how general ledger accounts are validated. General ledger accounts are validated when they are used in general ledger account fields.

Before you begin

You specify the system-level, default format of general ledger account codes using the GL Account Configuration window in the Database Configuration application. You specify organization-specific general ledger account codes using the Add/Modify Account Structure option in the Chart of Accounts application.

Procedure

1. Open the Chart of Accounts application.
2. In the Organizations table window, select the organization for which you want to specify validation rules.
3. From the **Select Action** menu, select **Validation Options**.
4. Select or clear the appropriate check boxes in the following list:

Option	Description
Deactivate GL Validations	If you deactivate general ledger validations, the entries in general ledger account fields are validated against values in the Chart of Accounts application. If you reactivate this setting, general ledger fields are not validated. You can enter values in general ledger fields, but cannot specify additional settings.
Validate GL Component Combinations	If you validate general ledger component combinations, only valid general ledger account entries are accepted. If you do not use this option, any combination of valid component values is accepted.

Option	Description
Validate Financial Periods	If you validate financial periods, checks are performed to ensure that a transaction occurs within an open, valid financial period. If you do not use this option, validations are not performed against defined financial periods.
Require Valid GL Account for All Transactions	If you require valid general ledger accounts for all transactions, transactions without a valid general ledger account are not allowed. Without this selection, valid general ledger debit and credit accounts must be present on all transactions.

5. Click **OK**.

Chapter 21. Working with cost management

You use the Cost Management application to track the financial resources required to complete the project. You can also link work orders from a project in the Cost Management application to the Work Order Tracking application.

Creating cost management projects

When you create a project in the Cost Management application you can assign work orders to that project in the Work Order Tracking application. You can then generate project cost information to track the financial resources required to complete the project.

About this task

After you assign the work order to a project or task, the project name appears in other applications in read-only format. To change the project or task to which the work order belongs, use the Work Order Tracking application.

Procedure

1. On the toolbar, click **New Project** and specify a project name.
2. Type a description for the project. By default, the application selects the **Is Chargeable** check box. This option allows you to charge costs to the project. If necessary, clear the **Is Chargeable** check box.
3. Specify a type for the project.
4. In the **Parent Project** field, assign the project to a parent project.
5. Optional: Type the **Budget** and **Budget Line** information.
6. In the **Value** field, specify the amount of money allocated for the project.
7. Specify the status of the project, the start date, and the end data for the project. By default, the system assigns an APPR status to the project.
8. In the Tasks table window, click **New Row** to add tasks to the project.
9. Click **Save Project**.

Chapter 22. Managing currency codes

You can define and manage currencies that you can use in purchase orders, purchase requisitions, invoices, and so on.

Creating currency codes

You create currency codes to define and manage currencies. A currency code record consists of a currency code, its description, and a setting specifying availability. All organizations can use currency codes.

About this task

You cannot use a currency that is not active.

Procedure

1. In the Currency Codes application, click **New Row**.
2. In the **Currency** field, provide a code value to represent the currency. For example, you use CND for the Canadian dollar.
3. Optional: Provide a description.
4. Optional: Deactivate the currency.
5. Click **Save Currency**.

Changing currency codes

You change a currency code record to change the currency code description or to make the currency active or inactive.

Procedure

1. In the Currency Codes application, find the currency code record that you want to change.
2. Activate or deactivate the currency.
3. Save your changes.

Chapter 23. Setting system properties

You can set system properties that various components use to control how applications and other aspects of the product work.

A system property is a key-value pair that is used at the system level.

Related reference:

Chapter 24, “System properties,” on page 419

You can modify the default values for system properties to tailor the system to your needs.

Global properties

A global property applies to the entire system. The property applies to all the server instances that use a common database, including a clustered environment.

Global properties always have the following options selected by default, but you can clear the options:

- **Online Changes Allowed** - By default, you can change most properties by using the System Properties application.
- **Live Refresh** - By default, most properties allow a live refresh where a new value applied to a system property takes effect immediately. **Live Refresh** is an action item available from the **Select Action** menu or from the application toolbar.

Instance properties

An instance property applies to a specific system server. When you create both a global value and an instance value for the same property, the instance value takes precedence.

Example

You can configure the **mxe.crontask.donotrun** system property to be an instance-specific property by specifying a specific server, such as MXServer1. You specify a value that applies only to that server, such as the bulletin board cron task, BBCron. As a result, BBCron does not run on MXServer1. However, BBCron can run on another server instance, such as MXServer2.

Options for system properties

Every property defined in the System Properties application has options that you can manage.

Table 64. System property option

Property options	Description
File override	Specifies whether the property and its value are loaded from a file rather than from the database.
Global only	Specifies whether this property must exist only at a system-wide level, and implies that the property cannot be overridden at the instance level.
Instance only	Specifies whether this property must be defined at the instance level. If the property must be defined at the instance level, you provide an instance-specific value and the property is not a global value.

Table 64. System property option (continued)

Property options	Description
Online changes allowed	Specifies whether the System Properties application is used to change the value for the property. For example, the global property mxe.db.driver does not allow online changes.
Live refresh	Specifies whether the property value can take effect immediately after saving the value.
Encrypted	Specifies whether the property is stored in an encrypted manner in the underlying product database. The value is encrypted by using the standard encryption functions for the product. For example, the global property mxe.int.uddipassword is encrypted.
Security level	Specifies the level of access to this property by various product components. Public The property and its value can be accessed through unauthenticated client sessions. Secure The property and its value can be accessed through authenticated client sessions. Private The property and its value can be accessed only with the business object framework of the system.
User defined	Specifies whether the property is created by a user or if the property is provided with the product.
Nulls allowed	Specifies whether the property can have null values. You can change this characteristic only for user-defined properties.
Data type	Specifies the type of value that can be provided for the property. The value can be an integer, alphanumeric, or a yes or no (YORN). For example, the global property mxe.allowLocalObjects is associated with the YORN data type. If you specify a value other than 1 or 0, an error message is displayed that indicates the value is invalid.
Domain	Specifies a domain that provides a list of values to which the property can be set. For example, the mxe.db.transaction_isolation global property is associated with the TRANSISO domain. Therefore, the values for the property must match a corresponding domain value.
Masked	Specifies whether the global and default values are hidden on the user interface.

System properties and encryption algorithms

The default encryption algorithm for system properties is DESede. You can configure different properties for the CRYPTO and CRYPTOX data types.

Table 65. Supported encryption algorithms.

Algorithm	Provider	Comments
AES	Cryptix, Sun	For Oracle Sun Microsystems, Inc., use mode = ECB.
Blowfish	BouncyCastle, Cryptix	
CAST5	Cryptix	
DES	Cryptix, Sun	
DESede	Cryptix, Sun	
IDEA	Cryptix	
MARS	Cryptix	
PBEWithMD5AndDES	Sun	For Oracle Sun Microsystems, Inc., you must use CBC and PKCS5Padding; key must be 8 bytes long.
PBEWithSHA1AndDES	BouncyCastle	
RC4	BouncyCastle, Cryptix	
RC6	Cryptix	
Rijndael	Cryptix	
RSA	BouncyCastle	Uses ECB and NoPadding (or empty string for mode and padding); spec is the private exponent, key is the public exponent.
Serpent	Cryptix	

Table 65. Supported encryption algorithms. (continued)

Algorithm	Provider	Comments
SKIPJACK	Cryptix	Spec length must be a multiple of 10.
Square	Cryptix	
Twofish	Cryptix	

Related concepts:

“System properties that contain password information”

Several system properties contain password information. Therefore, if you change a password, you must update the associated property value.

System properties that contain password information

Several system properties contain password information. Therefore, if you change a password, you must update the associated property value.

When you change a password, you must update the values of the following system properties:

- **mxe.adminPasswd**
- **mxe.adminusercredential**
- **mxe.b2b.password**
- **mxe.db.password**
- **mxe.int.uddipassword**
- **mxe.report.bo.rptServerLogonPass**
- **mxe.system.regpassword**

For example, if you change the database password, you must update the **mxe.db.password** property.

Related reference:

“System properties and encryption algorithms” on page 414

The default encryption algorithm for system properties is DESede. You can configure different properties for the CRYPTO and CRYPTOX data types.

“Security properties” on page 437

The data types Crypto and CryptoX are used to encrypt passwords and other types of confidential information. You use security properties to specify security levels for your organization, such as the data that must be encrypted and can be decrypted.

Values of system properties in files and applications

To simplify working in a development environment, you can assign a value to a system property in both a file and in the System Properties application. The value in the file takes precedence.

If the file override option is selected, the property is defined in the **maximo.properties** file.

In the System Properties application, if a system property is configured to be available only from the **maximo.properties** file, but the property is not present in this file, then the application server does not start and a message is written to the **maximo.log** file.

If a property is defined in the **maximo.properties** file, but not defined in the System Properties application, the property is not loaded at startup. A warning is written to the **maximo.log** file.

Example

The following are examples of the advantages of assigning a property value in both a file and in the System Properties application:

- Multiple developers can use a common database, and can run separate system instances with different property values.
- When you want one server in a cluster to handle a specific cron task, you can create a **maximo.properties** file specifically for that server instance.

Restoration of default values for system properties

In the System Properties application, you can restore the global default values and system default values for system properties.

Global default values

When you restore the global default values, the instance-specific rows from the database that are not flagged as instance-only in the MaxProp database are removed for the global properties that you select. In effect, the user-defined global values apply to all instances of the application server. The exception are properties that are required to exist at the instance level.

To restore the global default values, click **Restore Global Defaults**.

System default values

When you restore the system default values, the non-instance-only values in the MaxPropValue database table with the product defaults listed in the MaxProp table are restored for the property names that you select. In effect, all user-defined values are removed and restored. The exception are the instance-specific values that are required to exist.

To restore the system default values, click **Restore Maximo Defaults**.

Fetch stop limit memory errors

A Java virtual memory heap out-of-memory error is caused by an operation that fetches and constructs too many objects into one Maximo business object (MBO) set. This type of error can disconnect all users from the server. You can use the **mxe.db.fetch** system properties to configure the fetch limit, which can help prevent out-of-memory errors.

You use the fetch stop limit properties to set an upper limit on the number of objects that are fetched from the database and constructed on the server into a single set. You can set different limits for different types of objects. When the upper limit is reached, an exception is thrown so that the process is stopped. You can then either reduce the number of objects by using filters or change the operation to prevent the database from fetching too many objects.

The fetch stop limit is enabled by default with a limit of 5000 for all objects. If the fetch stop limit is reached and the exception is thrown during a user interface operation, the error message is typically shown to the user. Depending on how the

outer logic is implemented, this message might be wrapped by another message or not shown. However, the error is always logged on the server.

Scenarios where fetch stop limit errors occur

Large fetch counts that cause errors typically occur in two scenarios:

- When you apply an action to a large result set. The result set is retrieved without sufficient filtering. This error can occur when you mistakenly start a list page action or when a legitimate operation loads too many objects.
- Where an operation embedded in a user action retrieves a large set of objects and the data cannot be filtered.

Actions to take when fetch limit stop errors occur

Analyze the data and identify the operation that causes the error. You can try to correct the error by filtering the data or by splitting the target of the action into smaller batches. If the error cannot be corrected, the administrator must adjust the system in one or more of the following ways:

- Increase the fetch limit for the particular object if it is safe to use a higher limit. If you do not want other user operations to be capped by this high number, you can call the `MboSet.setLogLargFetchResultDisabled(true)` method on the MBO set.
- Correct the data to avoid processing such a large data set.
- Use discardable MBO sets to prevent all the objects from being fetched into memory.
- Move the operation into a cron task and schedule the cron task to run at a low-demand time or on another server.
- Modify the process so that it fetches fewer objects.

Related reference:

“Database properties” on page 423

You can use system properties to help manage the database.

Chapter 24. System properties

You can modify the default values for system properties to tailor the system to your needs.

Asset properties

You can use system properties to help manage assets.

Table 66. Asset properties

Property	Description	Default value
mxe.app.asset.delete AttributesAssetMovedBack	Deletes attributes when an asset is moved back to its original site. The default value is 0 which is no.	0
mxe.assettopology.depth	Represents the maximum depth that you can go in the asset topology.	5
mxe.assettopology.init	Indicates the initial depth for the asset topology.	2
mxe.assettopology.maxnodes	Specifies the maximum number of nodes that can be shown in the asset topology.	200
recon.engine. dataset.map.ASSET	Represents comma-separated data set names that can be reconciled against assets.	DEPLOYED ASSET
recon.engine. dataset.provider.ASSET	Represents the data set provider class for assets.	psdi.app.recontask.engine. dataset.AssetDataSet

Attached document properties

You can specify property values to control the use of attachments across the system.

Table 67. Attached document properties

Property	Description	Default value
mxe.doclink.doctypes.allowedFileExtensions	Represents the type of files that are allowed to be attached to files on the application server computer. The application server must be restarted for the changes to take effect.	pdf, zip, txt, doc, docx, dwg, gif, jpg, csv, xls,xlsx, ppt, xml, xsl, bmp, html
mxe.doclink.doctypes.topLevelPaths	Represents the top-level doclinks directory that stores all document folders on the application server computer. The Default File Path value entered in the Manage All Documents Folder window must be a subdirectory of this value. This value cannot be null. The application server must be restarted for the changes to take effect.	The default location is \DOCLINKS, typically under the default drive C:
mxe.doclink.path1 through mxe.doclink.path10	Specifies the HTTP server path to link documents that are attached to records. This property needs a live refresh only.	

Table 67. Attached document properties (continued)

Property	Description	Default value
mxe.doclink.doctypes.defpath	Represents the default path for the doclinks folder on the application server computer. This folder is where the physical documents that are attached to a record are stored. This value must be configured before you can use the system and needs a live refresh only.	
mxe.doclink.maxfilesize	Represents the maximum file size in megabytes for the doclinks folder that can be uploaded. You can change this property in the System Properties application; however, you must also rebuild the Enterprise Application Archive (EAR) file.	10
mxe.doclink.multilang.aix.websphere	Indicates whether the system is running on IBM WebSphere Application Server on AIX.	False
mxe.doclink.multilang.hpux.websphere	Indicates whether the system is running on IBM WebSphere Application Server on HP-UX.	False
mxe.doclink.multilang.linux.websphere	Indicates whether the system is running on IBM WebSphere Application Server on Linux.	False
mxe.doclink.multilang.solaris.websphere	Indicates whether the system is running on IBM WebSphere Application Server on Oracle Solaris.	False
mxe.doclink.multilang.windows.websphere	Indicates whether the system is running on IBM WebSphere Application Server on Microsoft Windows.	False
mxe.doclink.defaultPrintDocWithReport	Specifies whether printable attachments are printed by default when you print a report.	True
mxe.doclink.securedAttachment	Hides the file path of documents that were added by using the Attached Documents feature. You can use this property for security reasons. The hyperlink for the attachment shows the application content root and the encrypted file name only.	False
mxe.doclink.securedAttachmentDebug	Helps you to troubleshoot when you are using secured attachments to view the attachment file. When security for the attachments is active and you set the value for this property to true, the attached file name and file path display in the application server log.	False
proxy_hostname	Indicates the external host that enables direct printing.	
proxy_port	Indicates the external port that enables direct printing.	

Doclink path translations

10 properties for **doclink** path translations are provided. However, you can add properties to the MaxProp table by using property maintenance. The `LinkedDocumentInfo` class loads the properties that it finds. The `LinkedDocumentInfo` class reads **doclink** properties to get path translations.

You can specify the property value as the native operating system path + "=" + http translation. The `<PATH>` tag is used in the properties file, and because of problems in specifying the `:` character, the `:` character is permitted in the database.

The table provides an example of properties for **doclink** path translations.

Table 68. Example of properties for doclink path translations

Current format	New property name
C<PATH>\\Doclinks	mxe.doclink.path1
<i>install_home</i> /mxadmin/DOCLINKS	mxe.doclink.path2

The MaxPropValue table for **doclinks** contains the values listed in the table.

Table 69. Examples of values for doclink path translations

Property	Value
mxe.doclink.maxfilesize	10
mxe.doclink.doctypes.defpath	C:\DOCLINKS\
mxe.doclink.path1	C:\Doclinks=http://documentserver/
mxe.doclink.path2	<i>install_home</i> /mxadmin/DOCLINKS=http://documentserver/

Automation scripts properties

You can use system properties to help manage automation scripts.

Table 70. Automation script properties

Property	Description	Default value
mxe.script.attributelevel	Represents the maximum relationship depth for a launch point variable with ATTRIBUTE binding type.	3
mxe.script.drivers	Represents a comma-separated list of script drivers that is useful when any custom, none JSR223 compliant, script engine needs to be plugged into the script framework.	com.ibm.tivoli.maximo.script.JSR223ScriptDriver

Bidirectional language properties

You can use system properties to enable and manage bidirectional language support.

Table 71. Bidirectional language properties

Property	Description	Default value
<code>mxe.bidi.support.on</code>	Indicates whether bidirectional support is enabled. If the value is empty, illegal or 0, the bidirectional support code is not enabled. A value of 1 indicates that the bidirectional support is enabled.	
<code>mxe.bidi.text.direction</code>	When bidirectional support is enabled, this property indicates that the base text direction of text which appears on the GUI is enforced. The value which is enforced is stored in this property. This property is the system level property for base text direction. There is also a Maximo business object and a Maximo business object attribute level for the base text direction property. The values for this property are: <i>LTR</i> , <i>RTL</i> , or contextual values.	

Bulletin board property

You can use the property to track viewed and unviewed messages in the bulletin boards.

The bulletin board property is **PMBBISTRACKED**. This property is used to set the global variable for bulletin boards to track viewed or unviewed messages. The default value is 1.

Calendar property

You can use the property to specify the system base calendar.

The calendar property is **mxe.baseCalendar**. This property indicates the type of calendar that is defined as the system base calendar. The default value is gregorian.

Classification item properties

You can use classification items (CI) properties to specify data set provider classes for CI and actual CI.

Table 72. Classification item properties

Property	Description	Default value
<code>recon.engine.dataset.map.CI</code>	Represents comma-separated data set names that can be reconciled against CI.	ACTUAL CI
<code>recon.engine.dataset.provider.ACTUAL CI</code>	Represents a data set provider class for actual CI.	<code>psdi.app.recontask.engine.dataset.ActualCIDataSet</code>
<code>recon.engine.dataset.provider.CI</code>	Represents a data set provider class for CI.	<code>psdi.app.recontask.engine.dataset.CIDataSet</code>

Communication template property

You can use the property to specify whether bind variables are ignored in the communication templates.

The communication template property is **mxe.comm.ignoreunresolvedbindings**. This property ignores unresolved bind variables in the communication templates. The default value is 0.

Condition property

You can use the property to list comma-separated table names that are excluded when the reference for a condition is checked.

The condition property is **mxe.condition.excludeCheckReference**. This property represents comma-separated table names that are excluded when the reference for a condition is checked. There is no default value.

Cron task properties

You can use properties to manage cron tasks.

Table 73. Cron task properties

Property	Description	Default value
mxe.crontask.deleteBB	Removes expired bulletin board records.	0
mxe.crontask.donotrun	Use ALL to exclude all cron tasks from running. To exclude a specific cron task from running, specify the <i>crontaskname.instanceName</i> instance.	
mxe.crontask.historycleanuprate	Determines how often, in minutes, the excessive cron task history records are removed. There is no action if the value is set to 0.	180
mxe.cronTaskInitDelay	Represents the cron task monitor initialization delay in seconds. After the system server starts, this property determines the amount of time before the server initializes the cron task.	60
mxe.cronTaskMonitorInterval	Determines the time intervals at which the cron task manager monitors the statuses of the cron tasks.	60

Database properties

You can use system properties to help manage the database.

Table 74. Database properties

Property	Description	Default value
mxe.db.autocommit	Represents the autocommit mode used for the Write connections. This property can be either true or false. The default is false, and you cannot change the default value.	0
mxe.db.DB2jdbcCollection	Represents the DB2 Java database connectivity (JDBC) collection. The <i>NULLIDR1</i> value activates query optimization REOPT ONCE.	

Table 74. Database properties (continued)

Property	Description	Default value
mxe.db.DB2LD TextCaseInsensitiveSearch	When set to 1, the DB2 search on Long Description field is not case-sensitive. Might cause performance delay.	0
mxe.db.DB2sslConnection	Represents the secure socket layer (SSL) connection.	False
mxe.db. DB2sslTrustStoreLocation	Represents the DB2 SSL truststore location.	
mxe.db. DB2sslTrustStorePassword	Represents the DB2 SSL truststore password.	
mxe.db.disableservercursor	Used to disable the server cursor.	1
mxe.db.driver	Represents the database driver. This property must be defined in maximo.properties file.	com.microsoft.sqlserver. jdbc.SQLServerDriver
mxe.db.fetchResultLogLimit	Determines the typical or largest fetch count for every object. A stack trace is created in the log file of the application server logs every time a multiple of the fetch log limit is reached. For example, if you set the property to 1000, the stack trace is logged at the 1000th record, the 2000th record, and so on.	5000
mxe.db.fetchResultStopLimit	Represents the fetch stop limit used when checking that the fetch stop limit is enabled. The limit applies to all objects for which a specific fetch stop limit property is not specified. A value of -1 means that there is no limit.	5000
mxe.db.fetchsize	Represents the size of the database fetch.	40
mxe.db.fetchsizeuse	Represents the flag that indicates whether to use the fetch size.	1
mxe.db.fetchStopExclusion	Provides a comma-separated list of object names. If an object name is in the list, the fetch stop limit check is disabled for the object. If an object name is in the list and the same object is specified in an mxe.db.fetchResultStopLimit.OBJECTNAME property, the exclusion overrides the other property. The values are represented in a comma-separated object name list in the MAXOBJECT table.	
mxe.db.fetchStopLimitEnabled	Enables or disables the checking of the fetch stop limit. Use 0 to disable and use 1 to enable.	1
mxe.db.format.date	Determines the database date function. A value of none tells the system to pass through the date value. You cannot change this value.	

Table 74. Database properties (continued)

Property	Description	Default value
mxe.db.format.nullvalue	Represents the database-specific format of the null value function.	- For IBM DB2, the value is <i>COALESCE</i>, and you cannot change the default value. - For Oracle, the value is <i>NVL</i>, and you cannot change the default value. - For SQL Server, the value must be set to <i>ISNULL</i>.
mxe.db.format.time	Represents the database time function. A value of none indicates that the time value is passed through. You cannot change the default value.	
mxe.db.format.timestamp	Represents the database time stamp function. A value of none indicates that the time stamp value is passed through. You cannot change the default value.	
mxe.db.format.upper	Defines the database uppercase function for the system. You cannot change this value.	Upper
mxe.db.initialConnections	Represents the number of database connections that are created when the application server is started.	8
mxe.db.logCorrelationid	Disables database cursor sharing. You can add the correlation ID as an SQL comment. Use the property only in debug mode.	0
mxe.db.logSQLPlan	Represents the log execution plan for full table scans.	0
mxe.db.logSQLTimeLimit	Represents the log that contains the SQL operations that exceed the time limit in milliseconds.	1000
mxe.db.lookupMaxRow	Represents the maximum number of records queried from database for lookups.	1000
mxe.db.lookupMultiplier	Used in conjunction with the mxe.db.lookupMaxRow property to show items.	5
mxe.db.maxFreeConnections	Represents the maximum number of free database connections that are available in the connection pool.	8
mxe.db.minFreeConnections	Represents the minimum number of free database connections that are available in the connection pool.	5
mxe.db.MLQBELooseSearchW0Join	Used for Oracle only. Do not use outer join for search on a multiple language enabled field.	0
mxe.db.newConnectionCount	Represents the number of new connections to be created when the minimum free connections are available in the connection pool.	3
mxe.db.optionnum	Represents the size of the option.	1000
mxe.db.optionuse	Represents the flag that indicates whether to use the option.	1

Table 74. Database properties (continued)

Property	Description	Default value
mxe.db.password	Represents the native database password for the Maximo connection. This property must be defined in maximo.properties file.	XXXXXX
mxe.db.proxyauthentication.mode	Represents the Oracle proxy authentication mode. This mode is only valid when you use Oracle Proxy DataBase Manager.	The values for this property are: • 1 = <i>username</i> • 2 = <i>username + password</i> • 3 = <i>distinguished name (DN)</i> • 4 = <i>certificate</i>
mxe.db.QueryTimeout	Represents the amount of time in seconds before the SQL query times out and is stopped.	300
mxe.db.refcount	Represents the reference count for the connection log.	100
mxe.db.resultsettype	The constant indicates the type for a result set object with the following characteristics: • The cursor can move only forward. For example: TYPE_FORWARD_ONLY public static final int TYPE_FORWARD_ONLY • The result set type that is scrollable, but not sensitive to changes that other users make. For example: TYPE_SCROLL_INSENSITIVE public static final int TYPE_SCROLL_INSENSITIVE • The result set type that is scrollable and sensitive to changes that other users make. For example: TYPE_SCROLL_SENSITIVE public static final int TYPE_SCROLL_SENSITIVE	TYPE_FAST_FORWARD
mxe.db.retrydbconnection	Reconnects to the database when you start the application server.	0
mxe.db.rowcount	Represents the SQL Server row count value.	0
mxe.db.schemaowner	Indicates the database schema owner. This property must be defined in maximo.properties file.	DBO
mxe.db.sqlinjection	Indicates whether the SQL injection check is enabled.	1
mxe.db.sqlserverPrefetchRows	Represents the setting to reduce lock contention and is only for SQL Server. The optimal setting is 200 rows. Setting a value larger than 500 can degrade performance.	0
mxe.db.sqlTableScanExclude	Indicates tables that should not have an execution plan logged.	

Table 74. Database properties (continued)

Property	Description	Default value
mxe.db.systemdateformat	Represents the system date format.	- For IBM DB2, the value is current timestamp - For Oracle, the value is sysdate, and you cannot change the default value. - For SQL Server, the value is getdate.
mxe.db.transaction_isolation	The installation sets the value to: TRANSACTION_READ_COMMITTED. You cannot change this value.	TRANSACTION_READ_COMMITTED
mxe.db.updateWithoutRowstamp	Indicates that you can allow updates on tables that do not contain a rowstamp column.	0
mxe.db.url	Represents the database URL. This property must be defined in maximo.properties file.	jdbc:sqlserver: //qadb02.swg.usma.ibm.com:1433; databaseName=SQL2K8R2B; integratedSecurity=false;
mxe.db.user	Represents the native database user for a Maximo connection. This property must be defined in maximo.properties file.	maximo
mxe.db.UseSiteListInQuery	Represents the use literal list for site and organization restrictions.	0
mxe.dbmanager	References the Java class of the Maximo database manager. This property also requires that you specify the jdbc database connection string as the Oracle call interface (OCI) connection string, and that you make the OCI driver accessible to the system web component Java virtual machine (JVM).	The default value is <i>psdi.server.DBManager</i>. If you have an Oracle database that requires proxy authentication, set this property to <i>psdi.server.OracleProxyDBManager</i>.
mxe.dbwatchdog.adminemail	Represents the administrator email address used to send the database connection watchdog mail.	
mxe.dbwatchdog.logInterval	Represents the interval time in minutes between when each database connection watchdog log is created. The database connection watchdog log is written to the database connection logger.	10
mxe.dbwatchdog.mailinterval	Represents the interval in minutes before a watchdog email is sent to the administrator.	60

Related concepts:

“Fetch stop limit memory errors” on page 416

A Java virtual memory heap out-of-memory error is caused by an operation that fetches and constructs too many objects into one Maximo business object (MBO) set. This type of error can disconnect all users from the server. You can use the **mxe.db.fetch** system properties to configure the fetch limit, which can help prevent out-of-memory errors.

Deployed assets property

You can use the property to specify the data set provider class for deployed assets.

The deployed assets property is **recon.engine.dataset.provider.DEPLOYED ASSET**. This property represents the data set provider class for deployed assets. The default value is `psdi.app.recontask.engine.dataset.DPADataset`.

Email interaction system properties

The email interaction system properties define the behavior and characteristics of the emails.

Mail key properties

The **@@MAILKEY@@** string specifies the appropriate signature string in the subject line of emails. For each configuration of email interaction, at least one of the following properties is included in the **@@MAILKEY@@** property:

Property	Configuration type
mxe.mfmail.STSignatureSimple	Object status change with a simple email format
mxe.mfmail.STSignatureAdvanced	Object status change with an advanced email format
mxe.mfmail.WFSignatureSimple	Workflow assignment with a simple email format
mxe.mfmail.WFSignatureAdvanced	Workflow assignment with an advanced email format

Email body properties

The following system properties are specified in the body of the emails that are sent in email interaction:

Property	Description
mxe.mfmail.ValueListBegin	Marks the beginning of the prompt section
mxe.mfmail.ValueListEnd	Marks the end of the prompt section
mxe.mfmail.AssistMarker	Identifies the characters that precede the prompt numbers
mxe.mfmail.LineSize	Indicates the line size in the email. The default is 72.
mxe.mfmail.adminEmail	Indicates the email address of the administrator to notify when errors occur that require administrative action. This property is optional

SMTP property

The **mail.smtp.starttls.enable** property affects all features that send email such as password changes.

Property	Description
mail.smtp.starttls.enable	Enables the use of the STARTTLS command if it is supported by the SMTP server. The default value is false.

Email listener properties

You can use system properties to help manage email listener properties.

Table 75. Email listener properties

Property	Description	Default value
<code>mxe.listener.rfc822depth</code>	Represents the number of levels required to retrieve attachments from a message file.	3
<code>mxe.listener.rfc822extension</code>	Represents the file extension for the message file that is downloaded.	
<code>mxe.lsnr.validateperson</code>	Indicates whether a person needs to be validated in the Email Listener application.	1

Environment properties

You can use system properties to help manage the system environment.

Table 76. Environment properties

Property	Description	Default value
<code>mxe.allowLocalObjects</code>	In production environments, this property is set to true to improve system performance. You set this property to false for development work or for custom applications.	1
<code>mxe.enableConcurrentCheck</code>	Allows multiple logins on the same user account. By default it has a value of 0, if you want to change that value then the property must be defined in the <code>maximo.properties</code> file. You set this property to true (1) to prevent multiple logins on the same user account. Before you create users, set this property to 1.	0
<code>mxe.MLCacheLazyLoad</code>	Represents the multiple-language metadata cache which loads 1 object at a time. You set this property to 1 to load all objects simultaneously for 1 language.	1
<code>mxe.useAppServerSecurity</code>	Represents the security that your configuration uses. By default, the security for the system is used. You set this property to true if your configuration uses security provided by an application server.	0
<code>mxe.UserLicenseKey</code>	Represents the product enabler (license key) that is used during installation. If the product enabler changes, you must update the value of this property.	

E-signature properties

Electronic signatures confirm that a person who modifies a record is the same person that logged in to the system. You can use the properties to manage the e-signature feature for your organization.

Property values

The e-signature property is: **mxe.esig.defaultuserid**. When this property is set to true, the default e-signature login is the login ID. The default value for this property is true.

General ledger property

You can use the property to control the use of the upper function for the GL account in QBE.

The general ledger property is **mxe.upperGLValues**. This property controls the use of the upper function for the GL account in QBE. The default value is 0.

Inventory property

You can use the property to set the maximum inventory usage line limit in an inventory usage document.

The inventory property is **mxe.inventory.maxInvUseLineLimit**. This property represents the maximum inventory usage line limit in an inventory usage document. The default value is 250.

Issues and transfers property

You can use the property to open the Long Operation window when an INVISSUE is saved and the processing time is long.

The issues and transfers property is **mxe.app.inviissue.doLongOpOnSAVE**. This property opens the Long Operation window when an INVISSUE is saved. There are two values: 0 that opens the HourGlass or 1 that open the Long Operation window. The default value is 0.

Lightweight Directory Access Protocol integration properties

You can use system properties to manage Lightweight Directory Access Protocol (LDAP) integration.

Table 77. LDAP integration properties

Property	Description	Default value
mxe.allowLDAPUsers	Indicates whether LDAP users are allowed into the system if they do not have a user record.	0
mxe.ClientCountMinutes	Represents the interval in minutes for counting sessions.	15
mxe.LDAPMaxErrors	Represents the maximum number of errors for LDAP or virtual machine manager (VMM) synchronization.	1000
mxe.LDAPUserMgmt	Indicates whether LDAP owns user management when <code>mxe.userAppServerSecurity = 1</code> .	1

maximo.properties file

You must define system properties in the `maximo.properties` file to ensure that the application server starts. If you do not define these properties, an error message is written to the log file of the system or to the application server console.

Location of the properties file

The `maximo.properties` file is in the `<Product_root>\applications\maximo\properties` folder. The table lists the properties that you define in the `maximo.properties` file.

Table 78. *maximo.properties* file

Property	Description	Default value
<code>mxe.name</code>	Represents the application server that binds the application server object to the Remote Method Invocation (RMI) registry.	MXServer
<code>mxe.rmi.enabled</code>	Indicates whether RMI is enabled. This property must be defined in <code>maximo.properties</code> for you to disable the property.	1
<code>mxe.rmi.port</code>	Represents the RMI communication port. If set at zero, RMI uses any available port. You can select another available port number.	0
<code>mxe.db.user</code>	Represents the database user that the server uses to attach to the database server. For IBM DB2, the user must be an operating system user. For Oracle, the user must be the schema owner. For SQL Server, the user must have a system administrator role as defined through <code>sp_addsrvrolemember</code> . For example, <code>mxe.db.user = MAXIMO</code> .	For Oracle, the value is <code>maximo</code> .
<code>mxe.db.password</code>	Represents the password for the database user name.	
<code>mxe.db.schemaowner</code>	Represents the owner of the database schema.	For IBM DB2, the value is <code>Maximo</code> . For Oracle, the value is <code>Maximo</code> . For SQL Server, the value must be <code>dbo</code> .

Table 78. *maximo.properties* file (continued)

Property	Description	Default value
mxe.db.url	Represents the Java Database Connectivity (JDBC) URL of the database.	For IBM DB2, the value is: <code>mxe.db.url=jdbc:db2://localhost:50000/dbalias</code> , where <i>dbalias</i> is the database name. For Oracle, the value is: <code>mxe.db.url=jdbc:oracle:thin:@dbserver:1521:sid</code> , where <i>dbserver</i> is the server name of your database server, <i>1521</i> is your default Oracle port number, and <i>sid</i> is your Oracle system identifier. For SQL Server, the value is: server name, port number, database name defined as: <code>mxe.db.url=jdbc:inetdae7a:servername:1433?database=databasename&language=us_english&nowarnings=true</code> , in which <i>databasename</i> represents the database name, <i>servername</i> represents the name of the server, and <i>1433</i> represents the default SQL Server port number. The string mxe.db.url=jdbc:inetdae can be followed by either 7 (supports Unicode) or 7a (supports ASCII). Currently, only ASCII for SQL Server is supported.
mxe.db.driver	Represents the thin driver defined in mxe.db.driver .	For IBM DB2, the value is: mxe.db.driver=com.ibm.db2.jcc.DB2Driver For Oracle, the value is: mxe.db.driver=oracle.jdbc.driver.OracleDriver For SQL Server, the value is: mxe.db.driver=com.inet.tds.TdsDriver
mxe.enableConcurrentCheck	Allows multiple logins on the same user account. You set this property to true (1) to prevent multiple logins on the same user account. Before you create users, set this property to 1.	0

Migration Manager properties

You use Migration Manager to migrate configuration content from one product environment to another. You can use the properties to control the migration of the configuration content.

Table 79. *Migration Manager properties*

Property	Description	Default value
mxe.dm.autoapprovepkgdef	Represents the Migration Manager auto approval of package definitions.	0
mxe.dm.collvalidlevels	Represents the maximum number of levels searched to find related records.	5
mxe.dm.collvalidsrcexclude	Specifies the source object, attribute, and value combinations that cannot be traversed to.	<code>MAXINTOBJCOLSMAXINTOBJALIAS</code>
mxe.dm.collvalidtgtexclude	Specifies the target object and attribute combinations that cannot be traversed from during validation.	<code>MAXINTOBJECT.USEWITH, MAXRELATIONSHIP.CARDINALITY, CRONTASKINSTANCE.SCHEDULE</code>
mxe.dm.continueonerror	Continues the Migration Manager deployment when an error is found.	0
mxe.dm.dbserver	Represents the database server name for Oracle 9.x.	

Table 79. Migration Manager properties (continued)

Property	Description	Default value
mxe.dm.dmroot	Represents the name of the Migration Manager root folder that is on the application server computer. This folder stores Migration Manager package files. You must configure this property before using Migration Manager.	
mxe.dm.dmsessiontimeout	Represents the timeout value of the Migration Manager HTTP session in minutes. When a long-running Migration Manager task (such as package creation or package deployment) starts, the timeout value for the session of the user who is currently logged in changes to the value specified by this property.	120
mxe.dm.dmstagecommit	Specifies the commit interval for records that are inserted into the staging table in the target environment for the Migration Manager. The value is specified in the source environment, and Migration Manager uses the value when distributing a package from the source database to the target database.	1 The default value indicates that the commit interval occurs for every (1) record.
mxe.dm.importlimit	Represents the collection import limit.	100
mxe.dm.previewfreememorythreshold	Represents the memory threshold in percentage before the preview operation stops.	20

mxe.help properties

The **mxe.help** system properties connect the user interface to the information center. Some of the properties are used to construct the link that opens the information center. To ensure that information center is available, match the values in the **mxe.help** properties to the information center that you deploy.

Table 80. Information center properties

Property	Description	Default value
mxe.help.host	Represents the host name or IP address of the information center.	127.0.0.1
mxe.help.maximohelpink	Represents the top-level help link. The property must match the event values in the menus.xml file.	com.ibm.mam.doc, welcome.html
mxe.help.path	Represents the path that is inserted between the information center port and the topic when the link to the information center is constructed.	/help/
mxe.help.port	Represents the port of the information center.	9080
mxe.help.protocol	Represents the protocol of the information system (HTTP or HTTPS).	http

Table 80. Information center properties (continued)

Property	Description	Default value
mxe.help.viewsearchtiplink	Represents the plug-in name and file name linked to View Search Tips , which is available from the List tab of most applications	com.ibm.mbs.doc, mbs_common/ c_advanced_search_tips.html

Related concepts:

“Online help configuration” on page 38

There are different deployment options for online help running in an information center.

OSLC properties

You can use system properties to help manage Open Services for Lifecycle Collaboration (OSLC) application integration.

Table 81. OSLC properties

Property	Description	Default value
mxe.oslc.dfltversion	Represents the default OSLC version for an OSLC provider.	2
mxe.oslc.prcreateep	Represents the Provider Registry Create Endpoint.	
mxe.oslc.prqueryep	Represents the Provider Registry Query Endpoint	PROVIDERREGISTRY
mxe.oslc.webappurl	Represents the Provider Public URL.	http://localhost/maximo/oslc/

Reorder property

You can use the reorder property to manage the reorder time out periods for your organization.

The reorder property is **mxe.reorder.previewtimeout**. This property represents the reorder preview time out period in minutes. This property is similar to the web server session time out. The default value is 30 minutes.

Report integration properties

You can use report properties to control how reports are created and managed. You can identify all the report properties if you specify report as a filter value in the **Description** field.

Table 82. Report properties

Property	Description	Default value
mxe.activex	Determines if ActiveX Controls can be used for Direct Print (DP) and Direct Print with Attachments (DPA). Enables printing of attached documents that are Microsoft file types (such as .xls, .doc, .ppt). If you do not want to enable Active X Controls to print Microsoft documents, you must set the value to N (0), then Microsoft documents are not used with DPA.	1
mxe.directprint.inherited.attachments	Defines if inherited documents are enabled for printing by using the Direct Print with Attachments function. Used for direct print with attachments.	0
mxe.directprint.javaconsole.debug	Enables you to output to Java Console for troubleshooting.	0
mxe.directprint.printtime.waitinminutes	Indicates the maximum duration the current print process should wait before bailing out.	10
mxe.report.AttachDoc.validateURL	Determines if the URL of the attached document needs to be validated before printing the document from the V7 Server.	1
mxe.report.birt.disablequeuemanager	Defines whether the queue manager is enabled. If the queue manager is disabled (value set to 1), scheduled reports are not run on the server. Used for multi-server configurations, performance maintenance.	0
mxe.report.birt.maxconcurrentrun	Manages the number of immediate and scheduled Business Intelligence Reporting (BIRT) reports that can be run concurrently Used for performance maintenance.	5

Table 82. Report properties (continued)

Property	Description	Default value
<code>mxe.report.birt.queueidletimeseconds</code>	Frequency, in seconds, that the queue manager polls the queue for new report jobs. Used for performance maintenance.	60
<code>mxe.report.birt.viewerurl</code>	Represents the BIRT Viewer URL for clustered or separate report servers, for example: <code>http://myhost:myport/maximo/report</code> Used for BIRT Report Only Server (BROS) configuration	
<code>mxe.report.adhoc.editWithGroupAccess</code>	Enables you to edit an ad hoc report when any security group has access to the report.	0
<code>mxe.report.cognos.content.store.package.location</code>	Represents the content store folder where the Cognos package is published.	
<code>mxe.report.cognos.datasources</code>	Represents the Maximo data source name that is used by the Cognos adapter.	
<code>mxe.report.cognos.db.schemaName</code>	Represents the Maximo database schema name that is used by the Cognos adapter.	
<code>mxe.report.cognos.db.sql.name</code>	Represents the Maximo database name where the database type is SQL-Server.	
<code>mxe.report.cognos.db.type</code>	Represents the Maximo database type that is used by the Cognos adapter.	
<code>mxe.report.cognos.namespace</code>	Represents the Cognos namespace which holds information about users, security groups, and roles.	
<code>mxe.report.cognos.serverURL</code>	Represents the Cognos dispatcher/gateway URI. The property is used to launch Cognos reports or administration.	
<code>mxe.report.custom.rptServerLogonPass</code>	Represents the password used to log on to the external report server.	
<code>mxe.report.custom.serverURL</code>	Represents the URL for custom report applications.	
<code>mxe.report.reportsInAPage</code>	Defines the number of reports that display in the Report window.	5
<code>mxe.doclink.defaultPrintDocWithReport</code>	Represents the default value for printing an attached document with report if printable type. Used for direct print with attachments.	True

Additional report properties

Depending on your external reporting system, you might need to specify additional property values. You can add properties in the System Properties application.

Security properties

The data types **Crypto** and **CryptoX** are used to encrypt passwords and other types of confidential information. You use security properties to specify security levels for your organization, such as the data that must be encrypted and can be decrypted.

CRYPTO and CRYPTOX parameters

Parameters identified as **mxe.security.crypto** are for the **CRYPTO** maxtype. These parameters identify the attributes that can be encrypted and decrypted.

Parameters identified as **mxe.security.cryptox** are for the **CRYPTOX** maxtype. These parameters identify the attributes that can be encrypted, but not decrypted. These maxtypes have their own means of encryption, the parameters for which are defined in the properties file.

Table 83. Security properties

Property	Description	Default value
mxe.sec.adduser.maxsets	Represents the maximum number of concurrent sets allowed for user self registration.	20
mxe.sec.forgotpassword.maxsets	Represents the maximum number of concurrent sets allowed for a forgotten password.	20
mxe.sec.IPblock	Performs security checks related to IP blocking.	1
mxe.sec.IPblock.MatchBoth	Matches both the client host and the client address when you check for clients that are blocked.	1
mxe.sec.IPblock.num	Represents the maximum number of incorrect login attempts allowed per number of seconds.	50
mxe.sec.IPblock.sec	Represents the time in seconds required for the IP blocking limit check.	30
mxe.security.crypto.algorithm	Identifies the attributes that can be encrypted and decrypted. Algorithm is the basic type of encryption that is used. This property can override the algorithm default value DESed.	
mxe.security.crypto.key	Identifies the attributes that can be encrypted and decrypted. The length of this property must be a multiple of 24.	

Table 83. Security properties (continued)

Property	Description	Default value
mxe.security.crypto.mode	Identifies the attributes that can be encrypted and decrypted. The following mode components are valid: • Cipher Block Chaining Mode (CBC) as defined in <i>FIPS PUB 81</i>. • Cipher Feedback Mode (CFB) as defined in <i>FIPS PUB 81</i>. • Electronic Codebook Mode (ECB) as defined in <i>The National Institute of Standards and Technology (NIST) Federal Information Processing Standard (FIPS) PUB 81, DES Modes of Operation, U.S. Department of Commerce, Dec 1980</i>. • Output Feedback Mode (OFB) as defined in <i>FIPS PUB 81</i>. OFB must use NoPadding. • Propagating Cipher Block Chaining (PCBC) as defined by <i>Kerberos V4</i>.	
mxe.security.crypto.modulus	Identifies the attributes that can be encrypted and decrypted. Modulus is used only for the RSA algorithm.	
mxe.security.crypto.padding	Identifies the attributes that can be encrypted and decrypted. The following padding components are valid: • NoPadding - No padding. • PKCS5Padding - The padding scheme described in <i>RSA Laboratories, PKCS #5: Password-Based Encryption Standard, version 1.5, November 1993</i>.	
mxe.security.crypto.spec	Identifies the attributes that can be encrypted and decrypted. The length of this property must be a multiple of 8.	
mxe.security.cryptox.algorithm	Identify the attributes that can be encrypted, but not decrypted. Algorithm is the basic type of encryption that is used. This property can override the algorithm default value (DESede).	
mxe.security.cryptox.key	Identify the attributes that can be encrypted, but not decrypted. The length of this property must be a multiple of 24.	

Table 83. Security properties (continued)

Property	Description	Default value
mxe.security.cryptox.mode	Identify the attributes that can be encrypted, but not decrypted. The following mode components are valid: • Cipher Block Chaining Mode (CBC) as defined in <i>FIPS PUB 81</i>. • Cipher Feedback Mode (CFB) as defined in <i>FIPS PUB 81</i>. • Electronic Codebook Mode (ECB) as defined in <i>The National Institute of Standards and Technology (NIST) Federal Information Processing Standard (FIPS) PUB 81, DES Modes of Operation, U.S. Department of Commerce, Dec 1980</i>. • Output Feedback Mode (OFB) as defined in <i>FIPS PUB 81</i>. OFB must use NoPadding. • Propagating Cipher Block Chaining (PCBC) as defined by <i>Kerberos V4</i>.	
mxe.security.cryptox.modulus	Identify the attributes that can be encrypted, but not decrypted. Modulus is used only for the RSA algorithm.	
mxe.security.cryptox.padding	Identify the attributes that can be encrypted, but not decrypted. The following padding components are valid: • NoPadding - No padding. • PKCS5Padding - The padding scheme described in <i>RSA Laboratories, PKCS #5: Password-Based Encryption Standard, version 1.5, November 1993</i>.	
mxe.security.cryptox.spec	Identify the attributes that can be encrypted, but not decrypted. The length of this property must be a multiple of 8.	
mxe.security.provider	Represents the security provider which is obtained from the policy file. The security provider is usually <code>com.ibm.crypto.provider.IBMJCE</code>. To use a different provider, you can specify a value for this parameter.	

Related concepts:

“System properties that contain password information” on page 415

Several system properties contain password information. Therefore, if you change a password, you must update the associated property value.

Server properties

You can use the server properties to control how a server operates, such as specifying values for the administration user, for user registration, and email authentication.

Table 84. Server properties

Property	Description	Default value
mail.smtp.host	Represents the name of the host that runs the SMTP server. This name is needed for facilities that use email notifications, such as workflow notifications, email, and any error message notifications. Your network administrator can provide this address.	na.relay.ibm.com
mail.smtp.sendpartial	Indicates that partial emails are sent to valid email addresses.	1
mail.smtp.ssl.enable	Enables SSL over SMTP.	FALSE
mail.smtp.starttls.enable	Enables STARTTLS over SMTP.	FALSE
maximo_extended_host	Represents the Maximo extended host.	
maximo_extended_host_protocol	Represents the Maximo extended host protocol.	
mxe.adminEmail	Represents the email address that is used if the user does not specify an email address in the labor record. This value is requested during installation.	
mxe.adminmode.logoutmin	Represents the number of minutes that users must log out before the application server is placed in Admin mode. Admin mode is used to configure the database (including the application of structural changes).	5
mxe.adminmode.numsessions	Represents the number of administrative sessions that are allowed after the application server is placed in Admin mode.	5
mxe.adminPasswd	Represents the password for the administrative user.	maxadmin
mxe.adminPassword	Represents the password for the administrative user.	maxadmin, the system login ID of the administrative user.
mxe.adminusercredential	Represents the credential of the administrative user.	
mxe.adminuserid	Represents the administrative user. This user must have access to all sites. The server uses this property for administrative tasks and to run cron tasks.	maxadmin
mxe.adminuserloginid	Represents the system login ID for the administrative user.	maxadmin
mxe.com.port	Represents the com port.	

Table 84. Server properties (continued)

Property	Description	Default value
mxe.convertloginid	This property identifies whether the login id that the user entered must be converted to uppercase before it is validated. For conversion, set the value to 1.	0 - do not convert
mxe.email.charset	Represents the character set for email notifications that are sent. When this property is defined, the character set is used to encode the subject and the message when an email notification is sent. For non- English environments, you must set the UTF-8 character.	
mxe.email.content.type	used to set the content type for all communications.	text/html
mxe.hostname	Represents the name of the workstation and port that hosts the MXServer.	localhost:7001
mxe.maxsequencecheck	Indicates whether to check for multiple sequence values before using.	0
mxe.registry.bindcount	Represents the retry count for the system Remote Method Invocation (RMI) registry binding. When the system server starts and it runs into registry bind failures, the server tries to start for the number of attempts specified in this property.	100
mxe.registry.port	Represents the server RMI registry port that the server components uses.	1099
mxe.retainrecord	Retains the state of the application tab child tables during editing (less than retainreclimit).	1
mxe.retainrecordlimit	Represents the maximum number of records for which the table can retain its state. A higher number can affect performance.	200
mxe.smtp.connectiontimeout	Represents the socket connection timeout value in milliseconds.	180000
mxe.smtp.password	Defines the password for email authentication. This property must be used with mxe.smtp.user property.	Null - disables email authentication.
mxe.smtp.timeout	Represents the socket I/O timeout value in milliseconds.	180000
mxe.smtp.user	Defines the user ID for email authentication. This property must be used with mxe.smtp.password property.	Null - disables email authentication.
mxe.system.DomainFactoryNameProvider	Represents the class path that provides the DomainInfoFactory class names used when initializing domains.	psdi.mbo.DomainFactoryNameProvider
mxe.system.regpassword	Represents the user registration login password. This value is requested during installation.	maxreg

Table 84. Server properties (continued)

Property	Description	Default value
mxe.system.reguser	Represents the user registration login name to register a user. The user name that is specified must have authorization to create new users. This value is requested during installation.	maxreg
mxe.usermonitor. InactiveSessionTimeLimit	Represents the time an inactive session remains in the MAXSESSION table.	120
mxe.usermonitor.timeout	Represents the session timeout period, in minutes, on the server to clear the cache.	30
mxe.userrestrictionlrucachesize	Manages the number of entries in the LRU cache that stores the restriction entries.	1000
WAS.LTAURL	Represents the URL used to launch the Log and Trace Analyzer.	/ibm/action/launch?pageID=com.ibm.ac.lta.web.ui.LogAnalyzer&showNavArea=false
mxe.server.enableCSRFBlocking	Enables checks for the CSRF security token.	1
mxe.service.runlist	Represents comma-delimited services that must be run.	
mxe.sessiontoken.timeoutseconds	Represents a token-based session timeout in seconds. This token is used when the user ID is being authenticated on the report server. The session token timeout setting code is generic and currently only BIRT Reporting uses this function. Used for BIRT Report Only Server (BROS) configuration	180 seconds

Utilities for logging and testing

To help you manage system performance, there are utilities that you can use for testing and debugging purposes. When you are satisfied with your deployment, you can disable these logging utilities.

Utilities that track system performance

By default, the mbocount logging utility, the logSQLTimeLimit logging utility, and the fetchResultLogLimit logging utility are enabled in the properties file. These utilities enable you to track the following possible system performance issues while configuring an initial system deployment:

- Excessive use of business objects
- Slow execution of SQL statements
- High number of records returned in a query result

To disable the logging utilities, change the debugging properties described in the table.

Table 85. Debugging properties

Property	Description	Default value
mxe.mbocount	Displays the number of business objects that the server created. To disable, change the value of the property to 0.	1
mxe.db.logSQLTimeLimit	Represents the SQL statements that take longer than the specified time limit are logged. The time is measured in milliseconds. To disable, change the value of the property to 0.	1000
mxe.db.fetchResultLogLimit	When this property is enabled, a stack trace is printed in the log for every business object set that fetches beyond the set limit of rows. The stack trace log is also repeated for every multiple of such fetches. To disable, change the value of the property to 0.	1000 rows
mxe.db.logSQLPlan (Oracle only)	Setting this property to true logs the execution plan for all SQL statements that contain a full table scan. If you define mxe.db.sqlTableScanExclude , all tables, except for the ones you intentionally exclude, are logged. If you do not define mxe.db.sqlTableScanExclude , only the SQL statements that exceed the time limit set in mxe.dblogSQLTimeLimit are logged.	0
mxe.db.sqlTableScanExclude=ACTION,MAXROLE,SCCONFIG,MAXUSER (Oracle only)	You can define the table names that you want to exclude from the log. The table names must be uppercase. If you define mxe.db.sqlTableScanExclude , all tables, except for the ones that you list, are logged. If you do not define mxe.db.sqlTableScanExclude and you set mxe.db.logSQLPlan=true , only the SQL statements that exceed the time limit set in mxe.dblogSQLTimeLimit are logged.	0
mxe.logging.CorrelationEnabled	Represents the correlation ID that is logged if a percentage of the queue is specified in the logger layout.	1
mxe.logging.disableLoggingPropFile	Allows you to disable the logging.properties file for logging.	0
mxe.logging.rootfolder	Represents the default root folder to where Maximo generated log files are written.	

Work order generation property

You can use to use a property to manage the generation of work orders.

The work order generation property is **mxe.msgLogFile**. This property represents the log file for work order generation.

Workflow properties

You can use workflow properties to specify the email account of the administrator and to control status inheritance for work items.

Table 86. Workflow properties

Property	Description	Default value
mxe.workflow.admin	Represents the email account of the workflow administrator. The workflow is not impeded if you leave this value blank, but error messages that have communication templates associated with them might not be sent.	
mxe.app.wo.flowControlStatusInheritance	Allows status inheritance for work items that are under process flow control. You set this property to true (1) to ensure that the status is inherited from the parent work order to the child work order. The Flow Start and Flow Complete statuses are not inherited; they are INPRG and COMP, by default.	0

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.



Printed in USA