Tivoli.

IBM

IBM Maximo Asset Management
IBM Tivoli Asset Management for IT
IBM Tivoli Service Request Manager

Version 7.1

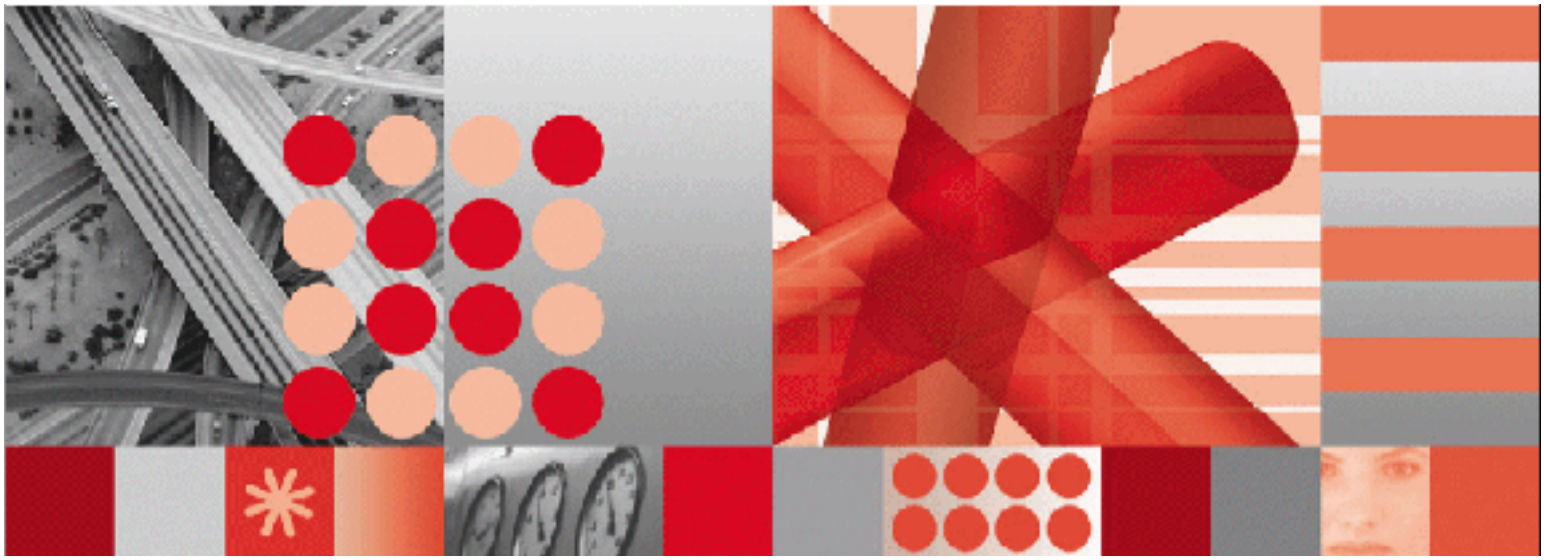

System Administrator Guide

Note
Before using this information and the product it supports, read the information in "Notices" on page 289.

This edition applies to version 7, release 1, modification 0 of IBM Maximo Asset Management, IBM Tivoli Asset Management for IT, and IBM Tivoli Service Request Manager and to all subsequent releases and modifications until otherwise indicated in new editions. This edition replaces any previous edition of this document.

Contents

About This Publication

This guide helps system administrators, network administrators, and database managers set up and configure the system, including managing the application server.

Some of the information in this guide may not pertain to your specific product.

Intended Audience

This guide is for database administrators, network administrators, and system administrators who perform the following tasks:

- Configuring and administering the application server
- Configuring and administering the applications
- Configuring and administering the database

1 System Overview

The system is based on service oriented architecture (SOA) and consists of several component servers. You access all applications using an internet browser (such as Microsoft® Internet Explorer®).

System Components

The system consists of multiple software servers. Depending on the size of your implementation, you can either install the servers on the same computer or on a separate server. Implementation consists of the following servers:

- Database server
- Application server

Application Server

The system is built using Java™ 2 Platform Enterprise Edition (J2EE) Technology, which requires a commercial application server. IBM WebSphere® Application Server is used. This server runs applications using JavaServer® Pages Technology, XML, and application-specific business components.

BEA® WebLogic® Server is an optional application server that you can use. However, users must provide this server themselves.

The user interface is rendered using XML, which lets you create common data formats and share the format and data. The XML code contains tags that reference each control in the user interface. The attribute values passed to controls in each XML tag determine the look and behavior of the controls.

The XML code is stored in the database, not within files. When accessing a system application, the application server loads the XML from the database. Then, based on the tags, the application server renders the user interface code sent to the client (Internet Explorer). Because the database stores the user interface data, any localizable text such as field labels, messages, and dialogs are also stored in the database.

Database Server

The following database servers are supported:

- IBM DB2® 8.2 FP7, or later, Enterprise Edition for Linux®, UNIX®, and Windows

- IBM DB2 9.1 FP2, or later, Enterprise Edition for Linux, UN IX, and Windows
- Oracle® 9i, Version 2
- Oracle 10, Releases 1 and 2
- Microsoft SQL Server 2005, Standard, or Enterprise Edition Service Pack (on Windows only)

System Requirements

System requirements depend on your operating system, database platform, and site configuration. See the *Installation Guide* for your product offering for minimum and recommended configurations.

Typical System Network Configuration

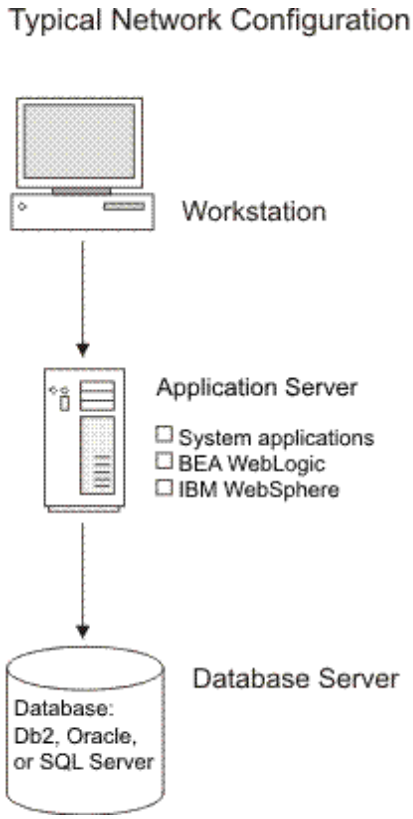

System, Organization, and Site Levels

System, Organization, and Site Level are terms that have special meaning. For example, a utility company owns several power plants, three water treatment plants, and two water distribution systems.

- System-Level refers to the entire company.
- The Organizations of a company are grouped into power plants, water treatment, and water distribution facilities.
- Each Organization has several Sites, which track inventory separately.

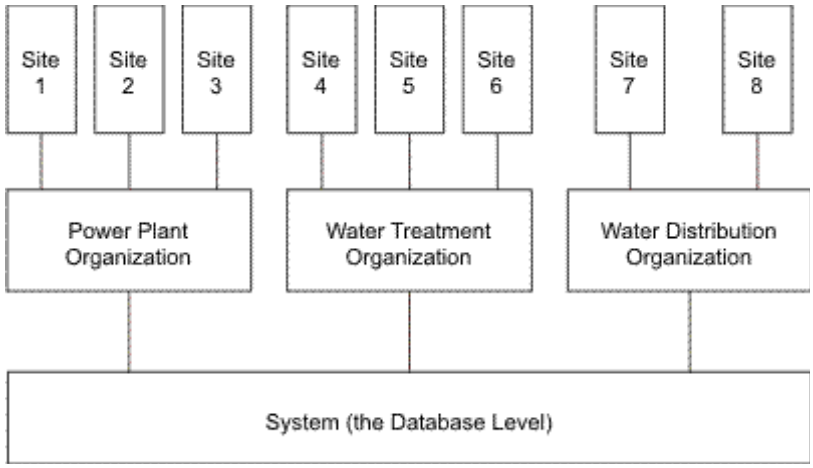

Creating System Settings

You create settings used throughout the applications. These settings include System, Organization, Site, and Security.

System Settings

You configure System-wide settings to create and activate Organizations, and must define at least one item in each of these steps.

Application	Description
Database Configuration	Provides options for configuring your Maximo database:

Application	Description
Sets	Create: ▼ Item Sets: groups of items that are shared between Organizations to enable things such as inventory sharing. In this part of the application, you name and define Sets. You do not add items here. ▼ Company Sets: groups of vendors that are shared between Organizations. In this part of the application, you name and define Sets. You do not add vendor companies.
Currency Codes	Define the currencies that you and your vendors use. Define codes and descriptions in this part of the system, and define exchange rates later, if applicable.
Organizations	Define Organizations and Sites. Many Organizations and Sites can share a single Maximo database. You must have at least one Organization and one Site to use the system. Autonumber Setup: This action lets you specify autonumber seeds and prefixes for record IDs that are unique at the System level.

Organization Settings

These applications let you configure Organization-wide settings and create Sites. Several of the following Organization options have defaults; check that the defaults correspond to your business rules.

Application	Function
Chart of Accounts	Define general ledger accounts and configuring rules surrounding general ledger account code validation. Most companies import their general ledger account codes and use this application to view the import and configure the validation rules. Configure additional Chart of Accounts options through Database Configuration.
Exchange Rates	Configure and administer exchange rates for currencies you and your vendors use.
Calendars	Define calendars, holidays, shifts, and work periods for your company. This data is used to schedule in other areas within the system.

Tab in the Organizations application	Function
Addresses	Configure addresses for your company.
Sites	Create Sites. You must have at least one Site to use the system. You can set additional administrative options for each Site in a different part of the application.

Action in the Organizations application	Function
Work Order Options	Configure the options that your Organization uses for work orders (examples: prompts for failure and downtime, rules for editing).
Inventory Options	Configure the options that your Organization uses for Inventory (examples: breakpoints, negative balances, reorder rules).
Drilldown Options	Configure the appearance of drill-down menus.
Safety Plan Options	Specify that the system displays the hazards of a work asset in the Select Hazards dialog box.
PO Options	Configure purchase order options for your Organization (example: how purchase requisitions are converted to purchase orders).
Contract Options	Associate terms and conditions with contract types.
Tax Options	Configure tax options for your Organization, including how multiple taxes are calculated.
PO Labor Options	Configure options for the rules of your organization for outside labor costs, including reporting of actuals and requirements for Purchase Orders.
Labor Options	Configure options for the rules of your organization for labor costs, including reporting of actuals and requirements for purchase orders.
Workflow Options	Configure options for your Organization's rules for Workflow processes, including reporting automated generation of work orders and purchase orders,
Autonumber Setup	Configure auto-numbering for items that are numbered at the Organization level, such as assets, to facilitate moves from one Site to another. For example, you can configure starting numbers and prefixes.

Site Settings

You can configure Site-level settings. Most options have defaults; check that the defaults correspond to your business rules. You control many additional options from individual applications.

Action in the Organizations application	Function
Work Order Options	Configure the Site-level settings for work orders (for example: numbering for tasks).
Inventory Options	Configure the Site-level settings for inventory for (example: how costs are calculated at issue).
PM Options	Configure how your Sites schedule Planned Maintenance (for example: scheduling by priority or frequency or how far in advance to generate work orders). The majority of system functions at the Site level, so you control most Site-specific configuration within individual applications, for example, work orders and purchasing.
SLA Options	Set SLA options.
E-Commerce Setup	If you implement the e-commerce Adapter, use e-commerce to configure vendor information for your Sites and the vendors with whom they do business.
Autonumber Setup	Configure auto-numbering for items that are numbered at the Site level, such as work orders. For example, you can configure starting numbers and prefixes. By default, Sites inherit auto-numbering from the parent Organization. You can change the default.

Security

The user community and their levels of access are defined below. For more information about groups and users settings, see Chapter 2, *Security* on page 11.

Groups A group defines levels of access to system applications and data.

Users A user record defines how the system looks and behaves for this user. The record must have an associated person record; you can maintain both records from the Users application.

A user must be a member of at least one group to access system applications. Users do not need to be given access to any sites in their groups to access data in System-level applications.

Reporting Options

Different reporting options are offered, including the following options:

- Business Intelligence Reporting Tool (BIRT)
- Business Objects®/Crystal Reports XI

- External Report Integration
- Open Database Platform Integration

Business Intelligence Reporting Tool

The system incorporates BIRT. The following report features are available in this tool:

- Use report templates and libraries to reduce report development time and ensure consistent user interface
- Select from a variety of report parameter options
- Schedule and send an e-mail message to reports at either a single point in the future, or on a recurring basis
- Drilldown to additional report details using hyperlinks
- Enable configurable toolbar access to minimize the clicks required to access reports
- Define record limits to maximize report performance
- Set up report security at either the application or individual report level

Business Objects/Crystal Reports XI

An integration to Business Objects Crystal Reports XI is available. This integration enables users to run Business Objects/Crystal Reports from within the system.

Report administrators can register and maintain the Business Objects/Crystal Reports from the Report Administration application. Users can then use the **Run Report** action to view the reports.

External Report Integration

The reporting function can integrate with many external reporting systems, including Oracle Business Intelligence, Hyperion®, Information Builders®, SAS®, and Cognos®.

Once the external report integration is enabled, users can view these reports from within the various system applications.

This function is report system and report version independent.

Open Database Platform Integration

The Open Database Platform Integration uses an external reporting system that derives information directly from the Maximo database, without integrating within the framework.

Because the system is based on an open database platform, this configuration can use any external reporting system to connect to the Maximo database. Once connected, users can create reports based on the information available in the system.

Understanding System Applications and Multisite

Applications store data at different levels in a multisite implementation. The following table lists which applications and functions are defined at each Multisite level:

Multisite Level	Applications or Functions
System or database — A System is a single instance of a Maximo database. A single System can contain one or more Sets, Organizations, and Sites.	▼ Attached Document Library ▼ Bulletin Board ▼ Classifications ▼ Collections ▼ Communication Templates ▼ Comparison Rules ▼ Computers ▼ Configuration Items ▼ CRON Task Setup ▼ Currency Codes ▼ Database Configuration ▼ Deployed Assets ▼ Domains ▼ E-Mail Listeners ▼ Escalations ▼ Features ▼ Incidents ▼ Job Plans ▼ KPI Manager ▼ Link Rules ▼ Master PM ▼ Meters ▼ Meter Groups ▼ Network Devices ▼ Network Printers ▼ Organizations ▼ People ▼ Person Groups ▼ Problems ▼ Reconciliation ▼ Relationships ▼ Report Administration ▼ Security Groups ▼ Service Level Agreements ▼ Service Requests ▼ Sets ▼ Solutions ▼ Ticket Templates ▼ Units of Measurement ▼ Users ▼ Workflow

Multisite Level	Applications or Functions
Set — Sets exist below the System level, but above the Organization level, to let multiple Organizations to share company and item data. Each Organization can have only one Company Set and one Item Set, but each Set can be shared by more than one Organization.	▼ Condition Codes ▼ Commodity Codes ▼ Company Masters ▼ Conversion values (Order Units/Issue Units) ▼ Item Master ▼ Service Groups ▼ Service Items ▼ Tools
Organization — An Organization identifies a unique legal entity. A large corporation can have different Organizations for different companies, or group all the facilities that exist in a continent or country into an Organization. There can be many Organizations in a single Maximo database.	▼ Calendars ▼ Chart of Accounts (General Ledger Account codes) ▼ Companies ▼ Crafts ▼ Exchange Rates ▼ Failure Codes ▼ Hazards ▼ Labor ▼ Labor Rate Contracts ▼ Lease/Rental Contracts ▼ Master Contracts ▼ Purchase Contracts ▼ Qualifications ▼ Software Contracts ▼ Software License View ▼ Tax Codes ▼ Terms and Conditions ▼ Warranty Contracts

Multisite Level	Applications or Functions
Site — A Site identifies a work location, such as a plant or facility.	▼ Activities ▼ Assets ▼ Assignment Manager ▼ Changes ▼ Condition Monitoring ▼ Cost Management ▼ Desktop Requisitions ▼ Inventory ▼ Invoices ▼ Issues and Transfers ▼ Labor Reporting ▼ Locations ▼ Lock Out/Tag Out ▼ Precautions ▼ Preventive Maintenance ▼ Purchase Orders ▼ Purchase Requisitions ▼ Quick Reporting ▼ Receiving ▼ Releases ▼ Request for Quotation ▼ Routes ▼ Safety Plans ▼ Stocked Tools ▼ Storerooms ▼ Work Order Tracking

For more information about multisites, see the *Multisite Administrator Guide*.

Security

2

Security Overview

Security systems use a two-step process: after authentication has verified that a user is who they say they are, authorization lets the user access various resources based on the identity of the user.

Authentication Overview

Authentication is the process of validating the identity of a user by providing proof that users are who they claim to be. There are different methods to authenticate users, and these methods share a common trait: authentication is always provided by a user ID and password. This process is distinct from authorization; authentication is not concerned with granting or denying access to system resources.

You can authenticate users through:

- Application server authentication can be used with an external authentication mechanism.
- Native authentication is provided to authenticate users and verify their identity and security authorizations.

Authorization Overview

Authorization determines which modules and applications a user can access, which actions they can perform, and which data they can view, modify, and delete. Authorization is provided by membership in one or more security groups.

Understanding Security Profiles

A user's security privileges controls a user's access to modules, applications, menu options, and data. All security access is based on security groups. A user name is associated with one or more security groups, and can have different levels of access to the applications and actions.

The Security Groups application is the building block for the security infrastructure. You configure security groups, either independent or non-independent, to provide narrow access or broad access to applications, Sites, labor, and other settings, such as general ledger components and approval limits and tolerances.

The security group to which a user belongs controls the user's level of access and privileges within the system. The system generates a user's security profile from all the groups in which a user is a member, using business rules to determine how the various security groups combine to build a virtual security profile. The security profile is like a fingerprint. It uniquely defines a user's access rights and privileges.

Creating security groups can be a modest to complicated task, depending on the number of sites in your company or facility, and depending on how fine-grained you want your security privileges.

Sites

The security architecture is designed to use sites as the first level of security for multisite implementations.

- If your system implementation has only one site, then, for each group, select the **Authorize Group for All Sites?** check box on the **Sites** tab.
- If your system implementation has multiple sites, create groups to represent each of the sites, all sites, or some logical grouping of sites within a security group (for example, a security group for site 1, and a security group for sites 2 and 3).
- Do not include any other privileges for the site groups.

If you select the **Independent of Other Groups** check box on the Groups tab, grant that group access to at least one site and one application, unless the group is used exclusively for system-level applications.

Applications, Storerooms, Labor, GL Components, Limits and Tolerance, and Restrictions

There are two strategies for the applications, storerooms, labor general ledger components, limits and tolerance, and restrictions:

- You can create groups that each 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 a user's profile as appropriate.
- You can create functional groups that combine some of the privileges. For example, you can create three different maintenance groups, each with differing levels of privileges for any or all of the properties covered by the different tabs. This strategy is good for defining groups in a detailed manner, so that when you associate one group with a user it 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.

Conditional Application Options

Another security feature is the ability to grant conditional access to applications. For example, you can define conditional access to an application that limits the user to use the Change Status action.

Viewing User Security Profiles

The Profile tab in the Users application displays the user's security profile. The security profile lists the settings and authorizations the user has after all of the groups are combined and merged. This information is read-only.

The security profile information displays the access that has been granted to the selected user through the combination of their roles. The information is displayed in a tree control using the standard Hierarchical Business Object interface. Because this data is stored in multiple tables, a virtual object is most likely required to implement the hierarchy.

Combining and Merging Security Groups

You can assign users to both independent groups and non-independent groups. Depending on the type of groups in which they are a member, the system generates a user's profile by combining, merging or combining, and merging a user's security groups.

The system can generate a security profile by combining the settings/sites of non-independent groups to the settings/sites of other non-independent groups. You can exclude a group from combining its settings/sites with other groups by making it an independent group. The settings for an independent group only apply to sites specified for that group.

The system can generate a security profile by merging the settings/sites for all independent groups. Finally, the system can generate a security profile by merging all independent groups with the settings/sites results set derived by combining all non-independent groups.

Security settings in the Users and Security Groups applications are at the system level, except for Approval Limits and Tolerances, which are organizational level settings.

When assigning applications to groups, it is important to understand whether the application is a system, set, organizational, or site-level application.

When you give users access to a system-level application, like Currency, it means that any changes the user makes in that application have a system-wide impact. For example, if you add EURO as a currency, it is available for all organizations and sites.

Similarly, if you change an application at the organizational level, that change applies to all sites in the organization. For example, if a Bedford user with Chart of Accounts access (an organizational level application) makes a structural change to an account, then that change affects all sites within the organization that Bedford belongs to, not just Bedford. Any changes that a user makes within a site-level application, like Assets, are limited to that site.

Finally, the level of the application controls the amount of data the user can see. For example, site-level applications display data for specific sites; organizational applications display data for all sites within an organization.

Rules for Combining and Merging Security Groups

When combining or merging groups to create a user's security profile, the system checks if there any rules that affect the selection or usage of security settings and authorizations.

Sites

The **Authorize Group for all Sites** check box on the Sites tab in the Security Groups application lets you give any members of this group access to all sites in the system. If a user does not have access to all sites using group membership, the system tracks to which unique sites a user has access through group membership, and the cumulative list of sites is included in the user's security profile.

Application Authorization

The application authorizations specified for an independent group apply exclusively to the sites or organizations associated with that group. However, the application authorizations specified for all non-independent groups apply to the totality of sites specified on all non-independent groups.

Access to the administrative applications, Users and Security Groups, is site independent when you do not specify a site. To modify site administration, you can, however, specify a site for the group that grants access to the administrative applications. For example, if a group is granted access to administrative applications at the Bedford site, a user who is logged into the system can only modify information for users who are associated with the Bedford site.

The accumulation of all unique application authorization records across security groups becomes the access list of Application Authorizations in the user's security profile.

The system populates the Action and Go To Menus with all of the options and applications that you granted a user, regardless of site or organization. You can grant a user the Change Status action for purchase orders in Bedford, but not Nashua.

In addition, if a user's security profile contains application authorizations but no sites, the user can access the applications but cannot view or insert records, except for the Users and Security Groups applications.

Finally, you can define conditional access to applications. For example, a user is a member of two groups, and one group has conditional access and the other group has unconditional access. The unconditional access overrides the conditional access, and the user has unconditional access.

Storeroom Authorization

The storeroom authorizations specified for an independent group apply exclusively to the sites associated with that group. In the Security Groups application, the Storerooms tab lets you authorize a group to make transactions with storerooms. The storeroom authorizations specified for all non-independent groups apply to the totality of sites specified on all non-independent groups. If you do not select the **Authorize Group for All Storerooms?** check box and do not authorize individual storerooms for the group, a user's security profile does not grant the user access to any storerooms at that site. However, if any of the groups to which a user belongs grants access to all or specific storerooms at a given site, then the user's security profile reflects the maximum amount of storeroom access. For example, if one group grants access to all storerooms at a given site, and a second group grants access to only one storeroom at the same site, the user's security profile grants access to all storerooms at that site.

TIP A user must have access to both a storeroom and the storeroom's site before the system adds the storeroom authorization to the user's security profile.

From an administrative standpoint, 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 has access to the administrative applications that enables the user to select the security group check boxes that authorize access to all storerooms. However, the user cannot add specific storerooms records using the **New Row** function.

The accumulation of all unique storeroom authorization records across security groups becomes the access list of Storeroom Authorizations in the user's security profile.

Labor Authorization

The Labor tab in the Security Groups application lets you authorize a group and any users who are members of the group for the following types of labor:

- 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 user's own labor records
- Individual labor records listed in the table window (LABORAUTH table)

If you do not select the **Authorize Group for All Labor?** check box and do not authorize individual labor options for the group, a user's security profile does not grant the user access to any labor. However, if any of the groups to which a user belongs grants access to all or specific labor options, the user's security profile reflects the maximum amount of labor access. For example, if one group grants access to all labor, and a second group grants access to only labor in the same crew as the user, then the user's security profile grants access to all labor.

TIP A user must have access to all labor or a subset of labor and a site in the labor's organization before the system adds the labor authorization to the user's security profile.

From an administrative standpoint, you can give a user access to the Users application and Security Groups application, but not grant that user access to any labor records. In this scenario, the user has access to the administrative applications that enables the user to select the security group check boxes that authorize access to all labor. However, the user cannot add specific labor records using the **New Row** function.

The accumulation of all unique labor authorizations across security groups becomes the complete list of Labor options available in the user's security profile.

Labor is an organizational-level application, which means that any user who is given access to Labor through a security group can view data for all sites in an organization, regardless of the site access granted by the group.

General Ledger Component Authorization

In the Security Groups application, the GL Components tab lets you authorize a group to change some or all of the general ledger components for sites and their organizations. To grant this authorization, check an individual component, such as Cost Center, or select the **Authorize Group to Change All GL Component Types?** check box to authorize the group to change all general ledger components.

If any of the groups to which a user belongs grants authorization to change all general ledger components or specific general ledger components, then the user's security profile reflects the maximum amount of general ledger component authorization. For example, if one group grants access to change all general ledger components, and a second group grants access to change only one general ledger component, then the user's security profile grants access to change all general ledger components.

If you do not select the **Authorize Group to Change All GL Components?** check box and you do not authorize individual components for the group, a user cannot change general ledger components.

The general ledger component authorizations specified for all non-independent groups apply to the totality of the applications, sites, and their organizations on all non-independent groups for the user. The general ledger component authorizations specified for an independent group apply exclusively to the applications, sites, and the organizations associated with that group.

The accumulation of all unique general ledger component authorizations across security groups becomes the complete list of general ledger component options available in the user's security profile. When users select general ledger components from the Select GL Account dialog box, they see the components they for which they are authorized in valid combinations only. When users use the Select GL Account dialog box within the Chart of Accounts application, the behavior is slightly different because the system presents all values for the authorized components so that they can create component combinations.

Approval Limits and Tolerances

In the Security Groups application, the Limits and Tolerances tab lets you specify limits and tolerances for a group. You can specify the following types of approval limits:

- Purchase Requisitions (PR)
- Purchase Orders (PO)
- Material Receipts (MR)
- Invoice
- Contract

The accumulation of all unique limits and tolerance authorizations across security groups becomes the complete list of limits and tolerance authorizations available in the user's security profile. The limits and tolerances that you specify for a group are at the organizational level, but users inherit authorizations for only those sites to which they have access.

If User 1 belongs to Group A with a purchase requisition limit of $5,000 and Group B with a purchase requisition limit of $10,000, then the user's security profile grants a purchase requisition limit of $10,000 for all sites to which the user has access, as long as the sites are in the same organization.

If the security profile gives User 1 access to two different organizations with different limits and tolerances in each organization, the user inherits the appropriate limits and tolerances for each site to which the user has access in each organization.

If Group A gives User 1 a $5,000 purchase requisition limit for sites in the EagleNA organization, and Group B gives User 1 a $10,000 purchase requisition limit for sites in the EagleSA organization, then the user has a $5,000 purchase requisition limit for all sites to which the user access in EagleNA and a $10,000 purchase requisition limit for all sites to which the user has access in EagleSA.

You can also specify the following types of tolerances for a group:

- Invoice
- Tax
- Service

For each type of tolerance, you can specify an upper dollar value and lower dollar value, and an upper percent and lower percent.

Example- Approval Limits

If two different approval limit values exist for the same limit, the system selects the higher value to apply to the user's security profile. For example, if Group A gives user Joe Black a $1,000 invoice approval limit and Group B gives him a $5,000 invoice approval limit, then Joe Black's security profile reflects the higher invoice approval, or $5,000. In some cases, the system must convert the values to the same currency using the EXCHANGE table to determine the higher value. This example is true for sites in Group A and Group B that are in the same organization.

Example- Invoice Amount Tolerance

If two different values exist for the same tolerance type, the system selects the higher value to apply to the user's security profile. For example, if Group A gives user Joe Black Invoice Amount tolerance and Group B gives him a $15 Upper Invoice Amount tolerance, then Joe Black's security profile reflects the higher Invoice Amount tolerance, or $15. In some cases, the system must convert the values to the same currency using the EXCHANGE table to determine the higher value. This example is true for sites in Group A and Group B that are in the same organization.

Example- Tolerances or Limits in Different Organizations

If two different values exist for the same tolerance type, but the groups that grant the tolerance amount have sites in different organizations, the system selects the higher value for sites within the same organization to apply to the user's security profile. For example, if Group A gives user Joe Black a $10 Upper Invoice Amount tolerance for sites in EagleNA and Group B gives him a $15 Upper Invoice Amount tolerance for sites in EagleSA, then Joe Black's security profile reflects two different Upper Invoice Amount tolerances: $10 for authorized sites in EagleNA and $15 for authorized sites in EagleSA. In some cases, the system must convert the values to the same currency using the EXCHANGE table to determine the higher value. This example is true for sites in Group A and Group B that are in different organizations.

Restrictions

Using restrictions, you can specify exactly which records are visible to members of a group. You do not use this feature to restrict access to applications and menus or to prevent a user from typing data; you control those privileges from the Applications tab.

If a user is a member of more than one non-independent group and has restricted access to one group, when you combine the groups, the user has the highest privilege.

For example, a user is a member of two non-independent groups with different levels of access:

- Managers group with access to pay rate information
- Maintenance group without access to pay rate information

The Managers group and Maintenance group are combined, giving the user the access to pay rate information in the Maintenance group.

A user within the EAGLENA and EAGLESA organizations is a member of the following two independent groups:

- Bedford Managers' in the EAGLE NA organization with access to pay rate information for laborers

- CHILEHDQ View in the EAGLESA organization with restricted access to the pay rate information

If the user accesses pay rate information in the Bedford Managers group, the records are viewable. However, if the same user accesses pay rate information in the CHILEHDQ View group, the records are not viewable.

Building Security Profiles

The size and complexity of your organization is an important factor in how you set up your security infrastructure using groups and profiles. For example, a simple implementation might only require that you build a few security groups for all the users at your company. A more complex implementation might require that you mix and match some security groups to create customized security profiles tailored to the specific needs of your organization. A global enterprise, for example, might require you to create an implementation team and a strategic plan to build security profiles that would meet the needs of users in multiple organizations located in different parts of the world.

Building Security Profiles: Examples

These examples show some approaches to construct security profiles for small to larger organizations.

Example 1

- Worker Group
- Management Group

Simple profiles constructed from a collection of security groups that contain sufficient application and Site access rights and privileges for all workers and management in the XYZ company.

The construction of the profiles shows how you can combine groups to restrict access to applications and limits and tolerances. For example, the worker security profile does not provide access to Purchase Requisitions and Financials, and does not provide a purchase limit.

Example 1 - Single Organization with Worker and Management Security Profiles

Example: Single organization with security groups that provide sufficient application, site and storeroom access and privileges for all users in XYZ company.

Worker and Management Groups for XYZ company

Worker Group

Sites	Inventory	Assets	Labor Reporting	Work Orders
✓ All sites	✓ Read ✓ Insert ✓ Save ✓ Delete ✓ All Actions	✓ Read ✓ Insert ✓ Save ✓ Delete ✓ All Actions	✓ Read ✓ Insert ✓ Save ✓ All Actions	✓ Read ✓ Insert ✓ Save ✓ All Actions
Storerooms ✓ All storerooms				

Management Group

Sites	Inventory	Work Orders	Labor Reporting
✓ All sites	✓ Read ✓ Insert ✓ Save ✓ Delete ✓ All Actions	✓ Read ✓ Insert ✓ Save ✓ All Actions	✓ Read ✓ Insert ✓ Save ✓ All Actions
Storerooms ✓ All storerooms			
Purchasing Limit ✓ $10,000	Assets ✓ Read ✓ Insert ✓ Save ✓ Delete ✓ All Actions	Purchase Requisitions ✓ Read ✓ Insert ✓ Save ✓ All Actions	Financials ✓ Read ✓ Insert ✓ Save ✓ Delete ✓ All Actions

All Workers →

Security Profile for all Workers

Application Access		Sites
Inventory Assets	Work Orders Labor Reporting	✓ All sites
✓ Read ✓ Insert ✓ Save ✓ Delete ✓ All Actions	✓ Read ✓ Insert ✓ Save ✓ All Actions	**Storerooms** ✓ All storerooms

All Management →

Security Profile for all Management

Application Access		Sites
Inventory Assets Financials	Work Orders Labor Reporting Purchase Requisitions	✓ All sites
✓ Read ✓ Insert ✓ Save ✓ Delete ✓ All Actions	✓ Read ✓ Insert ✓ Save ✓ All Actions	**Storerooms** ✓ All storerooms **Purchasing Limit** ✓ $10,000

Example 2

This sample Organization uses a mix of non-independent security groups, a type of group picklist, dedicated to individual security categories such as application access, Site access, storeroom access, and approval limits.

For example, the applications groups are for specific functional areas within the company, such as Maintenance and Purchasing.

When you add users to a related set of security groups, the system builds security profiles for these users who provide the access rights and privileges that they require to perform their job responsibilities within a specific functional area of the organization.

Example 2 - Single Organization with Mixed Security Groups

Example: Single organization with a mix of non-independent security groups dedicated to individual group categories like application, site and storeroom access. Security profiles reflect functional areas within the company, like Maintenance and Purchasing, as you add users to groups that provide the required access and privileges needed to perform specific job responsibilities.

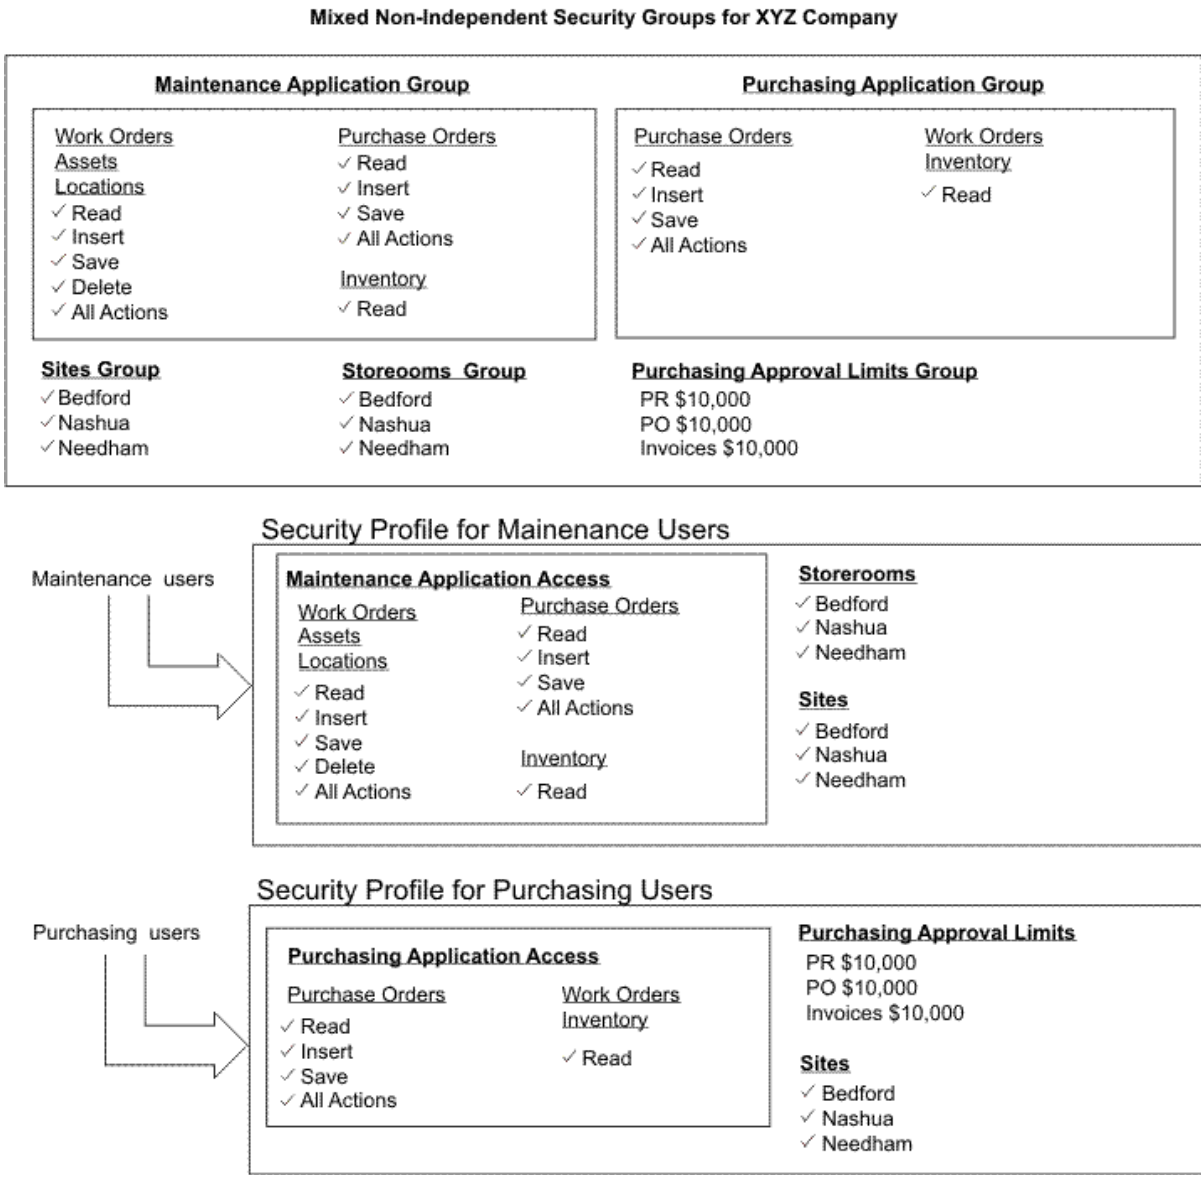

Example 3

This example contains a single Organization and independent security groups. The Organization uses Site administration so each Site group is limited to an individual Site with the Organization.

To provide read-only application access at a different Site for an employee, add the user to that Site group and to an application access group that provides read-only access to certain applications.

In this example, the user works in Nashua and has the most privileges and broadest range of application access. To provide access to Bedford and maintain Site administration, add the user to the Bedford Site group and the read-only application group.

Example 3 - Single Organizations with Independent Security Groups Using Site Administration

Example: Single organization that practices site administration with independent security groups. The application access groups are restricted to the Nashua and Bedford sites within the organization. The Nashua group provides the user with most of the application and storeroom access needed to perform his or her job. However, this user also requires Read-only access to several applications at the Bedford site. This example shows how to combine independent security groups so that a user has sufficient application access to perform his or her job responsibilities across sites.

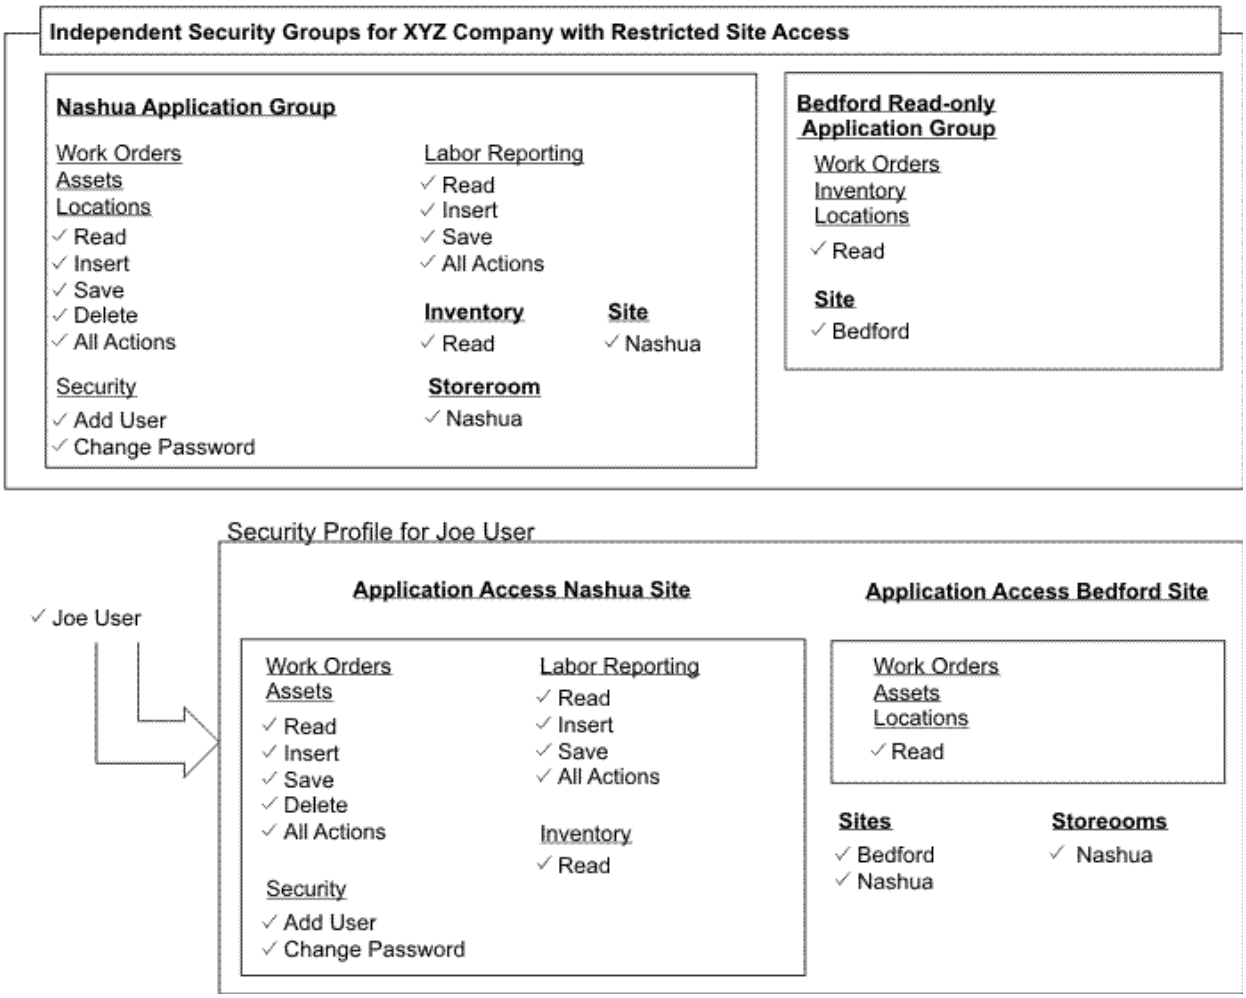

Example 4

This profile is built from independent and non-independent groups. The user's membership in one independent Site and read-only application group provides restricted, read-only access to several applications at remote Sites.

Membership in the non-independent groups provides sufficient application access rights and approval limits to perform the user's job responsibilities at the user's primary work Site.

Example 4 - One Organization with Independent and Non-Independent Security Groups

Example: Single organization that practices site administration with independent and non-independent security groups. The independent group provides the user with read-only application access at several remote sites. The non-independent groups provide the user with all the application access and approval limits he needs to perform his job responsibilities at his primary site.

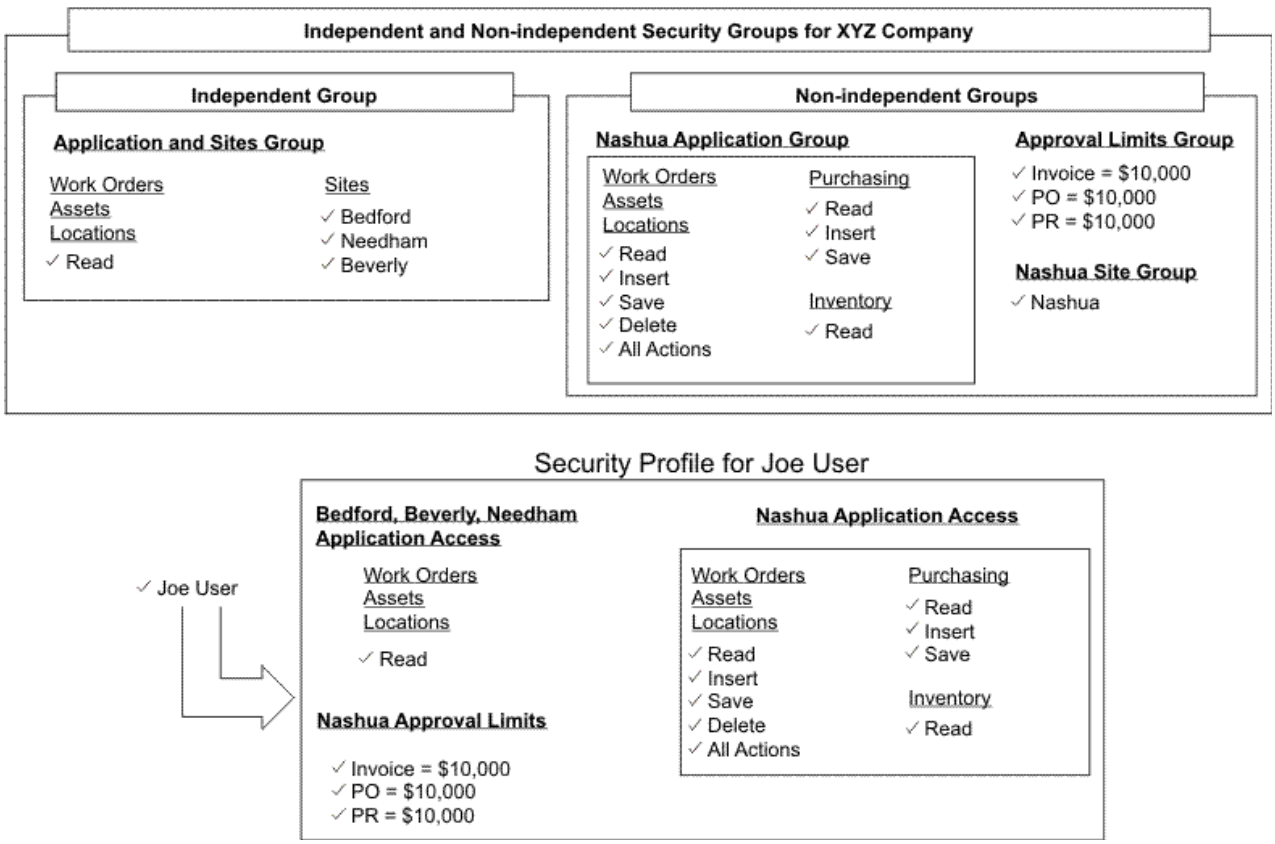

Example 5

This profile is for a user in a multi-Organization implementation (EagleNA and EagleSA) and is a member of these security groups:

- Three non-independent security groups, including a cross-Organization Site group, give the user the same application access and approval limits at a separate Sites within different Organizations (EagleNA/EagleSA)
- One independent Site and read-only application access group that provides application access for several Sites within a single Organization (EagleNA)

This example illustrates how user access to Organization-level and System-level applications applies to all Sites within an Organization and with the system.

Multi-Organization with Independent and Non-Independent Security Groups

Example: Multi-organizational implementation that uses independent and non-independent security groups to provide the user with read-only access to certain applications at several sites in one organization (EagleNA) and more robust application access at two other sites that reside in separate organizations (EagleNA and EagleSA). This example also includes some system and organizational-level applications, like Currency and Exchange Rates, that provide system-wide (all sites) and organization-wide (all sites in an organization) access.

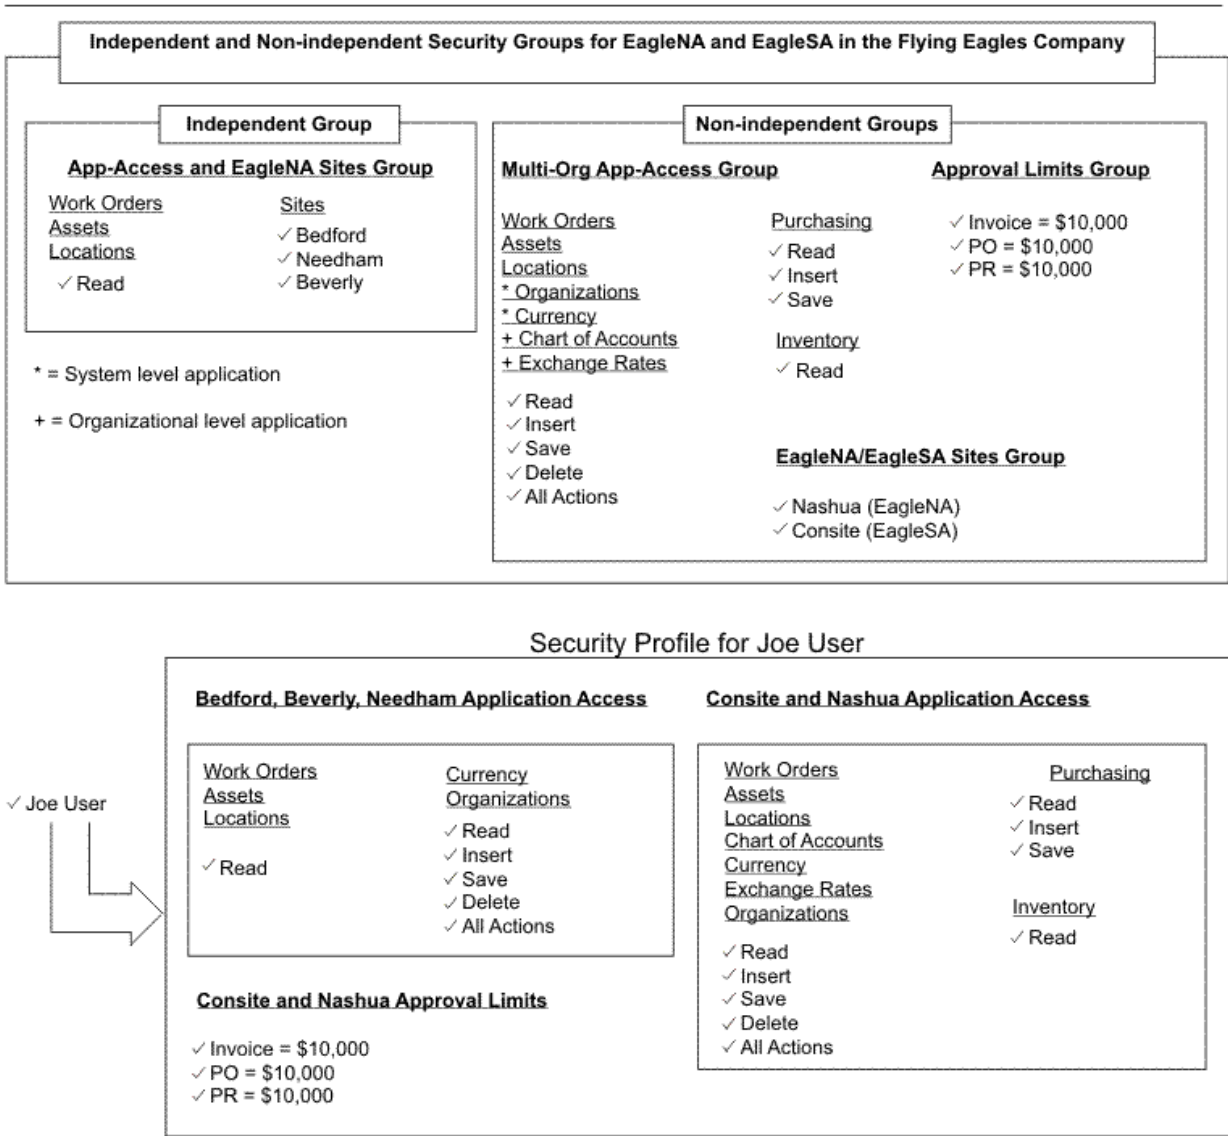

Example 6

This profile is for a user who is in a multi-Organization implementation (EagleNA and EagleSA) and who is a member of these security groups:

- Three non-independent security groups, including a cross-Organization Sitegroup, give the user the same application access and approval limits at a separate Sites within different Organizations (EagleNA/EagleSA)
- One independent Site and read-only application access group that provides application access for several Sites within a single Organization (EagleNA)
- The EVERYONE group as specified in the 'Group for all users' setting in the Security Controls dialog which is being used to configure global conditional access to an application option

This example shows how conditional access to an application option is reflected in a user's profile.

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.

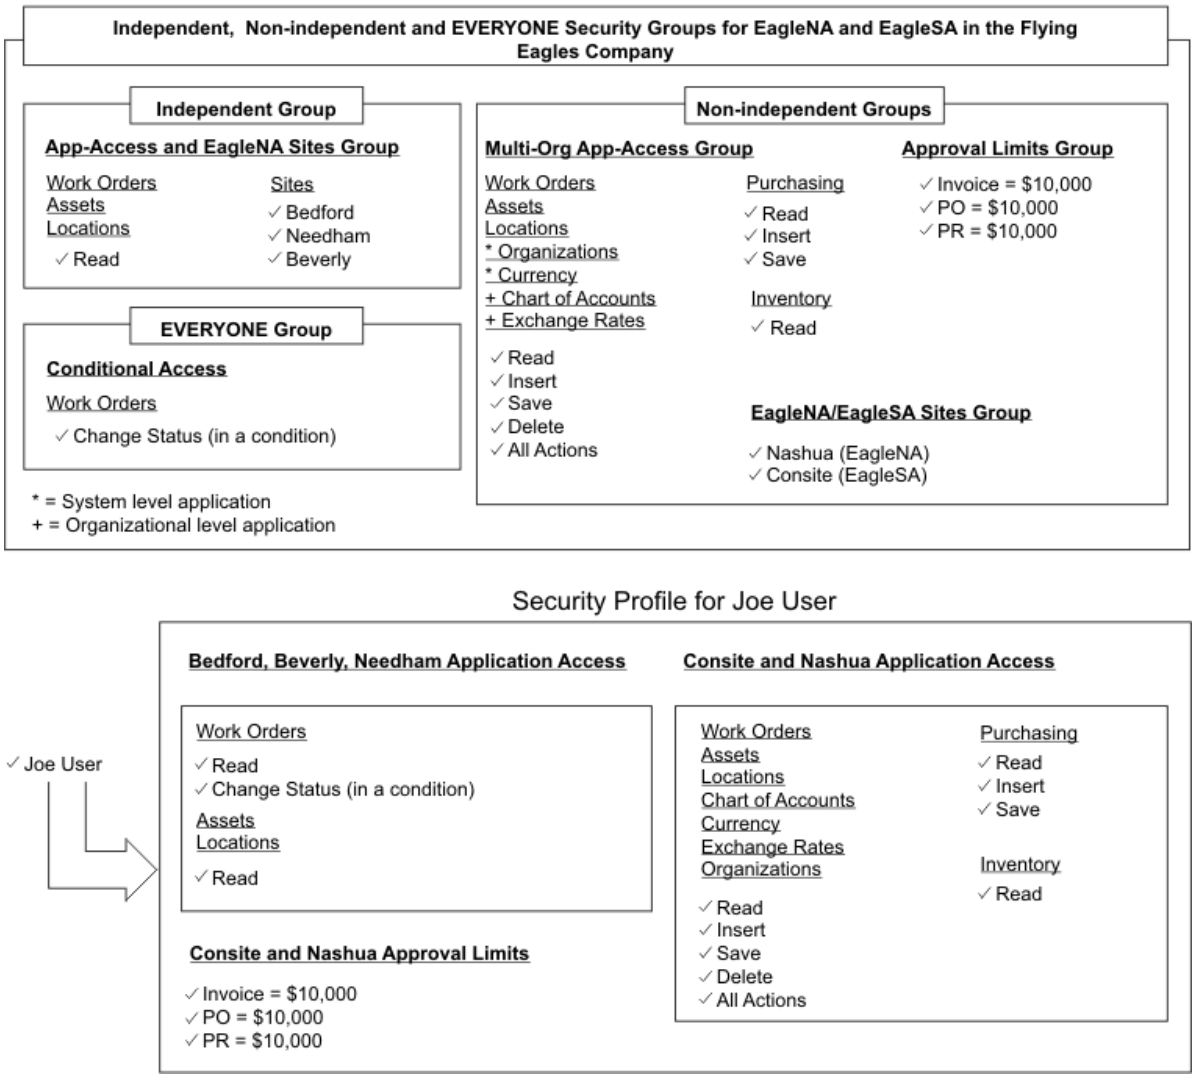

Setting Up Security Groups

The Security Groups application and Users application work together. Security groups define authorizations, and the users inherit the authorizations of the groups to which they belong, which then creates their security profiles.

For additional information, see *Authorization* on page 53.

Start Centers

Through portlets, the Start Center allows quick access to the tools and key performance indicators that users access typically. Administrators can configure the default Start Center that users sees when they access the system. To configure the Start Center, administrators must have access to the Start Center customization application within signature security.

Start Centers are assigned to security groups. Therefore, the first time user log into the system, they see a Start Center based on a template for the security group. If a user belongs to more than one security group, the user might see tabs at the top of the page, where each tab is a Start Center page for a different security group.

Administrators can grant users authorization to configure their Start Centers and control which portlets users see and can configure. The Start Center can contain the following portlets:

- Bulletin Board
- Favorite Applications
- Inbox/Assignments
- KPI Graph
- KPI List
- Quick Insert
- Result Set

For additional information, see *Start Center online help*.

Security Group Types

There are two types of security groups: independent and non-independent. When a security group is created, there is an **Independent of Other Groups?** option, which lets you specify if a group as independent or non-independent. If you do not specify a group as independent, the access rights and grants in this group are combined with access rights and grants in other groups that are also not independent.

These two types of security groups provide flexibility when you build the security infrastructure for your organization. Simple organizations might use one or two security groups; however, larger organizations with many users and a complex infrastructure might want to build some security groups that reflect varying levels of application and storeroom access and approval limits.

Understanding Independent Security Groups

When a security group is created, you specify a group type using the **Independent of Other Groups?** option. One option is to create an independent security group. An independent security group has the access rights and grants that cannot be combined with the rights and grants from other groups.

For example, a user is authorized to delete records at the Bedford site. For the Bedford site, the user would be granted the Delete Option action for the application and the site in the independent group.

Understanding Combined Security Groups

The concept of combined security groups and a derived security profile provides administrators with a powerful way to manage the security infrastructure within or across organizations.

The virtue of creating many groups is being able to combine them in many ways to fashion individual security profiles. A major attribute of a group is whether it is independent of other groups (this attribute is a check box on the Group tab). By default, this check box is clear, meaning that the group is non-independent and that you combine privileges when you combine groups. If you select the check box, the system does not combine privileges; the group is independent.

When you combine privileges, the highest privilege applies. If a user belongs to multiple groups that define the same privilege at different levels, the user has the highest privilege. For example, if group A has a purchase order limit of $5,000 and group B has a purchase order limit of $10,000, then a user who is a member of both group A and group B has a purchasing limit of $10,000.

Combining privileges becomes more useful as an implementation strategy when you have multiple sites. Typically, you set up groups that only define site access, for example, SITE1, SITE2, and SITE3. You define other groups to define application privileges, purchasing approval limits, and so forth. If you have a user for whom you created a security profile that includes SITE1 and some other groups to define application privileges, and you want the user to have the same privileges at SITE2, you add SITE2 to the user's profile. The user has the same rights in SITE2 and in SITE1.

On the other hand, you might want to define some groups as independent so that when you combine groups, a user has one set of privileges at one site and a different set of privileges at another site.

In practice, when you combine non-independent groups the resulting security profile gives members the greatest amount of access that results from combining the groups. However, if restrictions have been applied to the security group on the Data Restrictions tab, those restrictions are appended to a user's security profile and can reduce the access rights that were otherwise granted by the combined groups.

For example, if a user's security profile consists of two non-independent security groups, one group with full access to the People application and one group with restricted access that limits members to view only people whom they supervise, the resulting security profile would have full access. In short, when combining groups, full access overrides any group restrictions.

TIP Using the Security Controls action in the Security Groups application or Users application, you can specify the group for all users, EVERYONE, in the **Group for all Users** field. The EVERYONE group always combines, even if the group is specified as independent.

Applying Independent and Non-Independent Security Groups

When building a user's security profile, you can mix and match independent security groups or non-independent security groups in the following ways:

- Do not combine (all independent security groups)
- Combine (all non-independent security groups)
- Do not combine and combine (a mix of independent and non-independent groups)

The Security Groups application is the foundation of security infrastructure. You configure security groups, either independent or non-independent, to provide narrow access or broad access to applications, Sites, storerooms, labor, general ledger components, various approval limits and tolerances, and to set up data restrictions and users.

Setting Up Users

A user record must be associated with a person record. A user can associate an existing person with a user, or use the Users application to insert a new person ID.

In the Users application, you can add users and manage their security privileges. You can perform the following tasks:

- Add users
- Manage user passwords
- Manage user status
- Manage user sessions
- Assign users to security groups to create a security profile
- View user security profile
- Grant database access
- Specify various user defaults, such as default insert site, default storeroom, default language, and default GL accounts for purchasing
- Set system-wide security controls for password requirements, login tracking, and new user default groups
- Grant users the right to access inactive sites
- Specify which users can access a screen reader to assist in interacting with the system

In addition to granting rights to individual users, you can add, delete, and replace group privileges for multiple users at one time. You can also manage user status for multiple users.

Default Insert Site

When a user creates a site-level record, the **Site** field defaults to the value that is specified in the **Default Insert** field in the user record. When a user creates an organization-level record, the **Organization** field for that record defaults to the organization of the site that is specified in the **Default Insert** field in the record.

- Use the **Set Security Profile** action
- Specify the default insert site on a user record

In addition, a user can use the **My Profile** action to set their own default insert site

BEST PRACTICE A default insert site is not required. However, as a best practice, assign a default insert site. Without a default insert site, many applications will not function. For example, you will not be able to insert purchase orders or labor records in an application that is site level.

Filtering the Default Insert Site

When a new user logs into the system, the user can view records for all authorized sites. To change this setting, assign a user a default insert site to insert records. Once you assign a default insert site, the records that the user can view pertain only to the default insert site. You can also change the default insert site without reconfiguring the user.

Using the Profile navigation link in the navigation bar, user can change their assigned default insert site to a different site to which they have access.

For example, you are managing a group of users who have the following security settings:

- Access to the Assets application
- Access to multiple sites
- Bedford is the default insert site
- **Query Uses Default Insert Site** setting is enabled

When these users log into the Assets application, they can only see asset records from the Bedford site. In addition, if the **Query Uses Default Insert Site** setting is not enabled, these users can see asset records for all the sites to which they have been granted access in their security profiles.

The **Query Uses Default Insert Site** setting is a display filter that shows users only records from their default insert site. If the application is at the organizational level, however, this setting displays to users all records from all sites within the organization, not just the default insert site. If you clear this setting, users can view records for all the sites to which they have access.

Understanding User Statuses

A user record can have one of the following statuses:

- ACTIVE – Default status for new records. A user record must be ACTIVE to log in to the system.
- BLOCKED – User cannot log in to the system. An administrator can choose to block a user. If login tracking is enabled and the user types their user name or password incorrectly too many times, the system can also block a user.
- DELETED – User names cannot be reused. If you delete a user record, the user ID is retained in the database.
- INACTIVE – When a user is inactive, the user cannot log in to the system. Inactive user records do not appear in select value lists. A user record with a status of INACTIVE cannot be associated with new records.
- NEWREG – Default status for user records created by self-registration. This status is used to identify user records to route into a Workflow process.

When you add a user, their default status is ACTIVE.

Managing Users

When you insert new users, they are added to a default security group called DEFLTREG. You can set up and configure the DEFLTREG group with limited authorizations and privileges out of the box.

The **User Name** field, which defaults to the new user ID you typed when you created a user, is the login name the user uses when logging into the system. The user ID must be unique for all user records in the system. However, you can change the user name, which is case sensitive, to an employee number or e-mail address.

Additionally, when you add a new user, you specify the user's default insert site and storeroom. Most importantly, from a security standpoint, when adding or updating users, you can assign the user to one or more security groups. The combination of these security groups builds the user's security profile, which describes the set of authorizations and privileges a user has across sites and organizations.

The Security Profile tab shows the user's security profile after the system has combined all of the user's security groups. Sorted by site, the security profile is an expandable tree structure presentation of the user's virtual profile.

The system updates the following tables when you create a user:

Database Table Description

Database Table	Description
MAXUSER	Updated with new user data
PERSON	Updated with new person date (if the system needs to create a person for the new user)
PHONE	Stores all user phone information including home, work, cell, pager, and so on, and indicates which number to call first (primary)
EMAIL	Stores all user e-mail addresses including home, work, and alternates and indicates which address to e-mail first (primary)
GROUPUSER	Updates the relationship between security groups and the user
USERPURGL	Stores the default purchasing general ledger account for the new user (optional)
GRPREASSIGNAUTH	Stores the name of the user as a person with authority to add new users to a security group (optional)
MAXUSERSTATUS	Shows the user's current status such as Active or Blocked, and stores history of user status changes
PASSWORDHISTORY	Stores the user's current password, and stores history of password changes. Information in the Password column is encrypted, but other columns like USERID and DATE are not (if you enabled Password Duration)

Modifying Passwords for Default Users

If you modify a password for a default user (such as a system user for self-registration), then also modify the associated property. These properties include:

- mxe.adminPasswd
- mxe.system.regpassword

You can access these properties using the System Properties application.

Modifying a Database User Password

Use the **Change Password** action to change a database user password. You can set a database user password to match a system password; however, the database user password must support the password requirements.

Deleting Database Users

You can revoke a user's database access. When you select this action for a user with a value in the MAXUSER.DATABASEUSERID column, the current access to the system tables is displayed. As an administrator, you can modify these tables.

To delete a database user ID, click **Drop Database User**. This action lets an administrative user create, update, and revoke a user's database access. When you select this action for a user with a value in the MAXUSER.DATABASEUSERID column, the system displays the current access to system tables.

User Types

The system has user types 1 through 10. In the USERTYPE domain, update the descriptions to reflect the types of users in the license. The appropriate type must be associated with each user ID to ensure license compliance.

Understanding Administrative Users

Because of its high degree of flexibility and configurability, the system blurs the traditional line between administrators and users. You can grant any user access to any system or application function; there are no restrictions. From a security standpoint, administrative users have full or even restricted access to the Security Groups application and Users application, and the responsibility to implement and maintain security services, such as adding users, building profiles, or general site administration. Conversely, a regular user logs into the system and uses the applications.

Administrative users, regardless of the size of the client shop, might want access to one or more of the following applications:

Applications to Use for System Administration

- Organizations
- Calendars
- Classifications
- Sets
- Integration
- Database Configuration
- Domains
- System Properties
- Logging
- Communication Templates
- Roles
- Currency Codes
- Object Structures
- Launch in Context
- Cron Task Setup
- Application Designer
- Workflow Designer
- Workflow Administration
- Exchange Rates
- Chart of Accounts
- Users
- Security Groups
- Database Configuration
- Actions
- Escalations
- E-mail Listeners
- Web Services Library

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.

Understanding 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. You cannot delete a system user. MAXADMIN and MAXREG are system users that are part of the database shipped to customers. If necessary, you can select the **System Account?** check box in Users application to create new system users. Similarly, to delete a system user, clear the check box and click **Save**.

Understanding Database Users

When you create or update a database user ID, grant the user access to the system's tables. To grant access to the system's tables, click the **Object Name** icon and select a table object. Also, specify the level of access, either Read, Insert, Update, or Delete.

The system only creates database user IDs when you select this option. The system does not create operating system IDs for databases that require an operating system ID on the database server. If you implement databases with this requirement, create the operating system ID.

Access to tables is not assumed and must be granted. Rights are not defaulted out-of-the-box.

At the time of installation, if database users are created, some additional grants that the system uses must be supplied to the database user. The following commands detail the standard grants that the system requires:

- Create user maximo identified by the system
- Alter user maximo default tablespace maximo quota unlimited on maximo
- Alter user maximo temporary tablespace temp
- Grant create trigger to maximo
- Grant create session to maximo
- Grant create sequence to maximo
- Grant create synonym to maximo
- Grant create table to maximo
- Grant create view to maximo
- Grant create procedure to maximo
- Grant alter session to maximo
- Grant execute on ctxsys.ctx_ddl to maximo

In the Users application, to allow the system to give database access to users, the following commands detail the additional grants that are required for the system to create database users:

- Grant create user to MAXIMO
- Grant drop user to MAXIMO
- Grant create session to MAXIMO with ADMIN OPTION
- Grant alter user to MAXIMO

Understanding People, User, and Labor Records

When you create records for individuals, the system requires the creation of additional records in the following cases:

- Labor **-** You use the Labor application to create and manage labor records for employees and contractors who perform work on tickets and work orders. Labor records contain information about an individual's skills and qualifications. These records are used to plan and schedule work, and to track labor costs for tickets and work orders.

  A labor record must have a person record associated with it. A labor record needs a user record if the laborer is going to use the system to view work orders, report labor hours, and so forth. You can associate a labor and user with the same person record. As a best practice, you create craft records for different job skills and qualification records for certifications, and record that information about the labor record. Other resources records are optional, but not required. For example, a labor record can have:

  - One or more crafts
  - One or more skill levels associated with a craft
  - One or more qualifications

- Person - You use the People application to create and manage records for individuals. A person record contains basic information about an individual's name, address, contact information, and other generic information.

  A person record does not require any other resource records, such as craft, labor, user, and so forth. However, you must create a person record when you create a user record or labor record. You can associate a single person record with both a labor and user record.

  A person can be a user and a laborer or neither. For example, someone calling the service desk to make a service request does not need to be a user, but your company might require that a person record exist for that user.

  As a best practice, create a person record for any individual whose name appears anywhere on a record. For example, someone calling the service desk might not need a user ID to access the system. However, you can use person records to check if the individual is authorized to make a service request.

  To manage employee information, you can create person records for all of your employees and contractors. Alternatively, you can create person records for laborers and others who must access the applications as part of their jobs.

- User - You use the Users application to create and manage records for users. User records contain user names, passwords, and security profiles that determine which applications, options, and data a user can access. A user must have a person record. A user record can be associated with only one person record, and a person record can be associated with only one user record. You can associate a labor record and user record with the same person record.

Other resource records (labor, craft, and so forth) are optional, but not required. If you create new user records and you do not specify a value in the **Person** field, the system prompts you to create a matching person record for the user record.

You create a user record for anyone who must log into the system to view create or manage records.

Password Options

The **Security Controls** action is available from both the Security Groups application and the Users application. The **Security Controls** action allows you to specify the system-wide defaults for passwords that are described below.

Automatic Passwords

The automatic password generation setting allows you to specify that the system generates random passwords and sends an e-mail notification to users. In addition, you can specify the following settings:

- **Always E-mail Generated Passwords to Users (Never Display Screen)**

  The system sends an e-mail password notification to new users. If e-mail is not configured for the system, you cannot create a user.

- **Allow Generated Passwords to Be Displayed on Screen**

  Passwords can be sent through e-mail, but it is not required.

Password Requirements

You can specify the following password requirements:

- Minimum password length - The default minimum password length is six characters.
- The number of identical adjacent characters allowed in a password.
- Whether a password can contain a login ID.
- Whether a password must contain uppercase or lowercase characters, a number, or special character. These special characters are supported:
  - Ampersand &
  - Angle brackets < >
  - Asterisk *
  - At sign @
  - Back slash \
  - Braces { }
  - Brackets [ ]
  - Caret ^
  - Colon :
  - Dollar sign $
  - Equal sign =
  - Exclamation point !
  - Greater than sign >

- Hyphen -
- Less than sign <
- Number sign #
- Parenthesis ( )
- Percent %
- Period .
- Pipe |
- Plus sign +
- Question mark ?
- Semicolon ;
- Slash mark /
- Underscore _

- Whether the first and/or last character of a password can be a number. Numeric characters are: 1 2 3 4 5 6 7 8 9 0.

- Whether the first and/or last character of a password can be a special character.

Excluded Password List

The **Security Controls** action lets you create and manage a list of passwords that are excluded from use.

If implementation uses an application server to authenticate with an external directory (through LDAP), you cannot use the system to perform some functions. These functions include:

- Self registration - This function is not supported with an external directory.

- Setting or changing passwords and password hints - All password-related functions are managed by the directory.

You can also set the following functions:

- Duration of password, in days. If you do not want passwords to expire, leave the Password Lasts this Number of Days field empty.

- Advance warning of password expiration - To have the system notify users X days before password expiration, type a value in the **Days Before Password Expires to Warn User** field.

- Days before previously used password can be reused - If you do not want the system to check for password reuse, type 0 in the **Days Before Previously Used Password Can Be Used Again** field.

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/or 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.

Understanding the Self-Registration Process (Disabled with LDAP)

New users can click **register now** to self-register from the Welcome page. The self-registration process lets new users register with a minimum amount of information. The system assigns self-registered users to a default security group, but nothing else until their registration request has been routed to an administrator through workflow (if enabled) and is approved or rejected. Once approved, the administrator assigns the self-registered user to the appropriate security groups and notifies the user that they can use the system. Rejected users are told to contact their supervisor for assistance.

The name of the default security group for self-registered users is DEFLTREG. However, an administrative user can designate and configure any security group to replace the default group in the Security Groups application or Users application by updating the **Default Group for New Users** field in Security Controls. For example, you can configure the access rights and privileges of the default group for self-registered users to reflect your company's business rules.

After an administrator approves a new user, the user is instructed to complete the registration process by typing more information in the My Profile application. After submitting a self-registration request, the system creates a person record and user record for the user.

The following table shows the required and optional information a user provides to create a self-registration request.

User Self Registration Information and Optional Fields

User Information	Optional
First Name	Password Hint Question
Last Name	Answer
User Name	Supervisor
Password	Default Insert Site
Confirm Password	Default Storeroom
Primary E-mail	Primary Phone
	Language
	Locale
	Time Zone
	Additional Information

The system also defaults additional registration information which is not shown to the user.

User Self-Registration Hidden Fields

Hidden Field	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 `Y` for new users created through the Users application or through self-registration, and when and administrator changes a user's password
Query with Site	Defaults to `Y`
Person Status	Defaults to `Active`
Transaction Notifications	Defaults to `Never`
Workflow Notifications	Defaults to `Process`
Accepting Workflow E-mail	Defaults to `Y`

Configuring for User Self-Registration

You can configure the system to allow users to self-register. To enable self-registration, make the following configurations:

- If desired, in either the Users application or Security Groups application, use Security Controls to rename the **Default Group for New Users** (default value is DEFLTREG) and **Initial Self-Registered User Status** (default value is NEWREG).

- If you want users to have a basic set of authorizations when they self-register, in the Security Groups application, add the desired authorizations to the DEFLTREG security group (or to the new group name that you chose). By default, the authorizations for this group are limited to changing and expired password, and accessing the Start Center (but not any content on the Start Center).

- When a user submits a self-registration form, it is placed in the SELFREG workflow process. By default, the process notifies an administrator that a new registration is pending, and the administrator must approve the new user before the user can access the system. Modify this process to be specific to your organization. For example, you can change this process to approve self-registered users.

- When a self-registration is being processed, a series of e-mail notifications are available as communication templates. These notifications are sent as the registration progresses. You must modify the content of these templates to be specific to your organization. The templates include:

    - NEWSELFREG
    - REGNOTIFY
    - REGAPPROVE
    - SELFREGREJ

When a user's self-registration is approved, the user can log in to the system with the authorizations that you granted to the DEFLTREG security group (or to the new group name you chose)

Creating a User Record Using Self-Registration

Using the Security Controls dialog box, you can determine a default user status that specifies whether self-registered new users can access the system immediately (ACTIVE), or whether new users must wait for their registration to be processed (NEWREG, INACTIVE).

You can also access the Security Controls dialog box from the Users application.

Creating a Self-Registration Workflow Process

In the Security Groups application, you can configure the Users application to use Workflow to process self-registered users. If you enable Workflow, the following standard Workflow options are available:

- Start/Continue Workflow
- Stop Workflow
- View Workflow History
- View Workflow Map

For more information about workflow options, see the *Workflow Implementation Guide*.

Enabling the Self-Registration Workflow Process

After you configured the workflow process, enable and activate the workflow process for self-registration to function. To complete this process, select the appropriate boxes within the Workflow Designer application.

Setting Defaults for New Users

You can type any security group in the **New User Group** field to be the default registration group for new users. Out of the box, DEFLTREG is the default registration group. By default, the system gives users who self-register a NEWREG security status. The security status you assign to self-registered users determines whether the system enables workflow for new, self-registered users.

In the New User Defaults section of the Security Controls dialog box, you can specify a self-registration user status:

NEWREG	User cannot log in; workflow is enabled
ACTIVE	User can log in; workflow is disabled
INACTIVE	User cannot log in; workflow is disabled

For example, you can decide to set the default user self-registration status to ACTIVE so that self-registered users can immediately use the system. You cannot, however, assign the status NEWREG to new users that you create through the Users application. The system reserves the NEWREG status for users that you create through self-registration.

Administrative users cannot select the BLOCKED or DELETED status from the self-registered user status value list.

Users placed in the new user default group acquire the application access and authorizations configured for that group.

As an administrative user, you can create new users, but you do not have the authority to delete DEFLTREG from a user's security profile. To have delete authorization for users within a group, open your user record and add the group to your Authorize Group Reassignment list. The system does not grant you authorization to delete DEFLTREG because it is the default group for self-registered users as defined by NEWUSERGROUP in the system variable table (MAXVAR).

Setting Up Authentication

Planning User Authentication

There are two supported mechanisms to perform user authentication:

- Using a Web client for native authentication
- Using the application server and a Lightweight Directory Access Protocol (LDAP) directory server for authentication

Authenticating Users Against LDAP through Virtual Member Management

You can authenticate users against LDAP using Windows® Server Active Directory. If your organization has Virtual Member Management in place, consider using it to perform your authentication.

When you configure the application server to authenticate against an Active Directory, you create and manage users in the LDAP directory server. The VMM cron task updates the Maximo® 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 Maximo database because these records might 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 might require programming to configure.

Both BEA® WebLogic Server® and IBM® WebSphere® Application Server support authentication against Windows Server Active Directory.

Native Authentication

You can use the native authentication provided with the system to authenticate users and verify their identity and security authorizations.

A user types a login ID and password in the Welcome page. The security functions validate whether the user ID and password exist 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, when the security services load at system startup, they perform the following actions:

- Verify if Login ID is blocked or inactive
- Authenticate Login ID and updates password history (if configured)
- Establish user's default insert Site, Organization, and Person ID
- Establish the user's language, locale, time zone, and Start Center ID
- Route any Workflow assignments to the user's inbox (if Workflow processes are enabled)

The following procedure is the most common way to authenticate application access:

1. At the Web client login screen, users type a login ID (in the **User Name** field) and password.
2. Security services validate users' credentials against the Maximo database. This validation uses Java™ encryption to check the user in the Maximo schema/database.
3. The system checks users' security profiles. Based on the authorizations that they contain, the system grants users access to the applications.

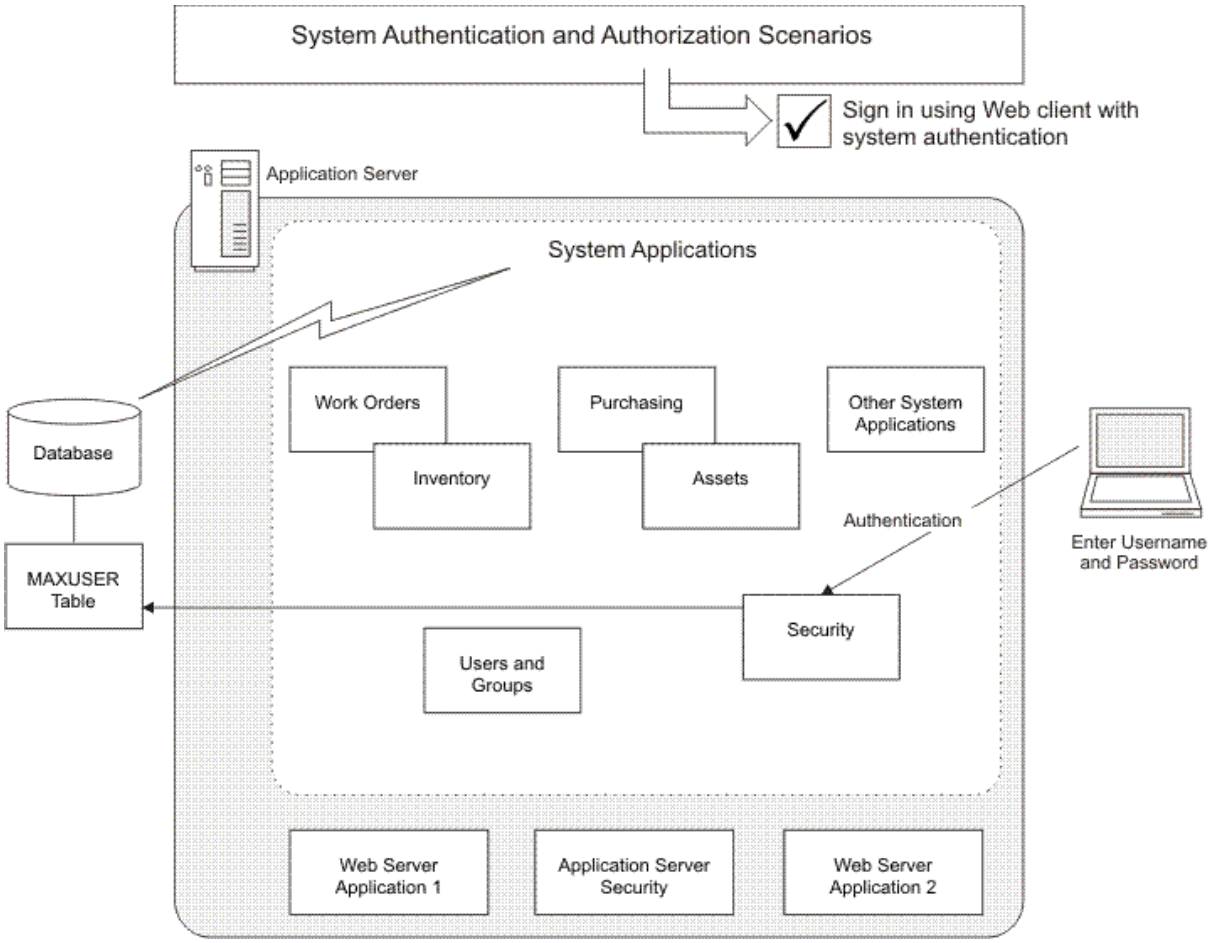

Security services load at system startup, and perform these actions:

- Verify Login ID (blocked or inactive).
- Authenticate Login ID and updates password history (if configured).
- Establish user's default insert Site, Organization, and Person ID.
- Establish the user's language, locale, time zone, and Start Center ID.
- Look for any Workflow assignments in the user's inbox (if Workflow processes are enabled).

Application Server Authentication

You can use application server security with an external authentication mechanism, such as LDAP to authenticate users. LDAP is a set of protocols to access information directories.

The system uses application server security with an external authentication mechanism. System installation enables WebSphere Application Server security by default.

Depending on the product offering, application server security might be enabled by default.

The system is built using J2EE technology, which requires a commercial application server. The system uses WebSphere Application Server or WebLogic Server.

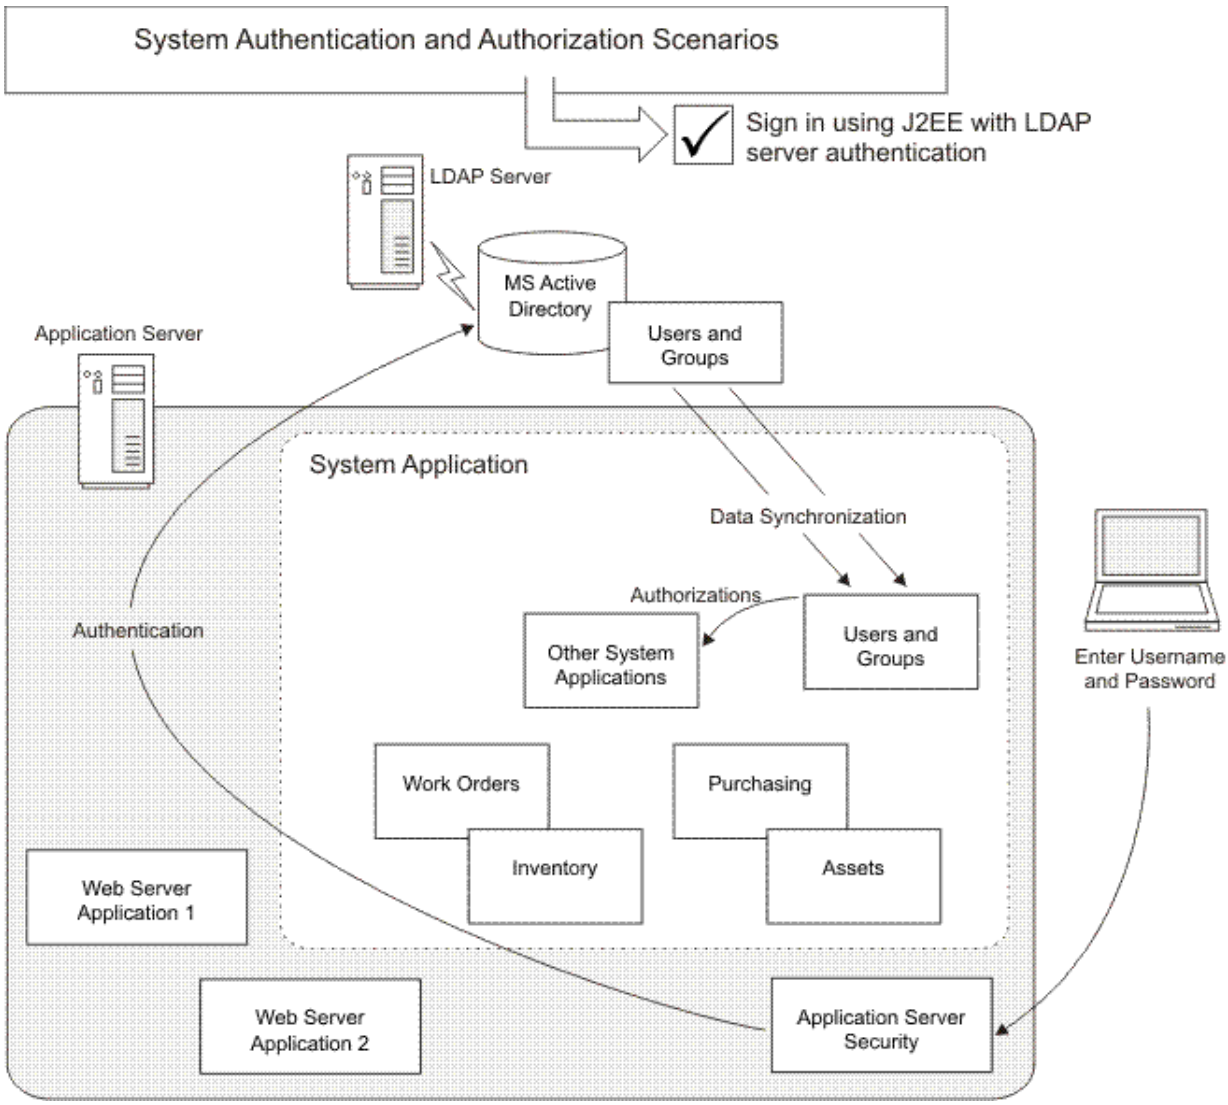

Application Server Security

Before configuring the system for WebSphere Application Server or WebLogic Server security:

- Review the following considerations
- Review the following preliminary tasks
- Configure the application server security

Considerations

- The application server is configured to authenticate against an LDAP server user registry.

  The system supports integration with Microsoft Active Directory or IBM Tivoli Directory Server LDAP server, and lets you move LDAP server data into Maximo database tables.

  The application server that you use determines which directory provides support.

- Add users and delete users and groups from the LDAP server, but provide authorization from the system.

  By default, the property mxe.LDAPGroupMgmt is set so that the directory owns group creation and group membership management.

- Configure all application-specific authorization rules for users and groups using the system's security module applications. In the system, disable:

- Password Information in Start Center

- Change Password application

- Self-Registration

- Users application

- Security Groups are limited to administrators assigning authorizations restrictions

- LDAP server users and groups are moved into Maximo database tables to identify users as system users, and provide user details in system applications.

  Users and groups deleted from the LDAP server are not deleted from Maximo database tables; audits might exist for users or groups.

- If user accounts are disabled from LDAP server, it takes several minutes for the application server to expire the user's cached information. The time it takes for cached information to expire depends on cache expiration settings.

- Before users can access the system:
  - The application server authentication must be passed.
  - Users' identities must exist in Maximo database tables.
  - The login ID is mapped to MAXUSER Login ID field.
  - Users must have authorization for at least one application.
- Application servers use Roles to identify users and groups with access to the system. All roles configured in an application are mapped to users or groups using application server-specific deployment descriptors or application server-provided administrative tools.

  By default, the system includes a security role, maximouser, mapped to the maximousers group, which identifies users with access to the system. You can change this role mapping to any user or group in the LDAP server.

Preliminary Tasks

These tasks fall outside the system environment. Before configuring the system to use WebSphere or WebLogic Server, and LDAP security, you must:

- Create the User Directory on an LDAP server.
- Create an organizational unit for the system.
- Create a group called maximousers, under the organizational unit.
- Create administrative users in the directory, and assign them to an organizational unit. The system requires these administrative users:
  - MAXADMIN
  - MAXREG
  - MXINTADM
- Assign these administrative users to the maximousers group.

Synchronization

Synchronization keeps data in the system current with data in the LDAP directory server. How data synchronization occurs depends on whether you use WebSphere Application Server or WebLogic Server:

- WebSphere Application Server - Data synchronization between LDAP repositories and the system is governed by the federated repositories which, in turn, are managed by Virtual Member Manager and the VMMSYNC cron task.
- WebLogic Server - Data synchronization between LDAP repositories and the system is managed by the LDAPSYNC cron task.

Synchronization data moves in one direction only — from the LDAP directory server to the system.

To synchronize data, create users and groups in the LDAP directory server. When you synchronize the users and groups that you created in the LDAP directory server with the system, the users and groups become users and security groups.

In the following procedure, you set up a cron task to synchronize data:

1. Log into the system as an administrative user.
2. Open the Cron Task Setup application.
3. Complete one of the following steps:
   a. If you are using WebLogic server, set the LDAPSYNC cron task to **Active**.
   b. If you are using WebSphere Application Server, set the VMMSYNC cron task to **Active**.
4. Set a schedule.

   Let the cron task run to synchronize all the users and groups from the LDAP directory server into the Maximo database tables.

Data Mappings

Set up the following parameters for the LDAPSYNC cron task and the VMMSYNC cron task.

LDAPSYNC Parameter	Description
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 SSL enabled
SynchAdapter	LDAP synchronization adapter
SynchClass	LDAP synchronization Class
SynchParameter	Parameter name, value pairs delimited by comma
UserMapping	The USER XML that the LDAP task uses

VMMSYNC Parameter	Description
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
SynchClass	VMM synchronization Class
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 the required columns must contain data. If all the required column data cannot be obtained from the LDAP directory server, specify default values.

To fill in default values for columns, the value must be enclosed inside {} brackets; for example, {ABC} fills in 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} and to generate system date use the notation {:sysdate}.

Synchronization Tips

These tips apply when you synchronize data.

TIP Information moves in only one direction — from the LDAP directory server to the system.

Deleting Users and Groups

Deleting users and security groups on the LDAP directory server does not delete them in the system. This restriction is for audit purposes for clients in regulated industries.

Disabling User Accounts

You must disable a user account in the system. If you disable a user account in the LDAP directory server, that user account is not disabled in the system.

Disabled users cannot log into the system.

If Synchronization Fails

If fields do not synchronize the LDAP directory server to the system, check the LDAPSYNC cron task parameters and VMMSYNC cron task parameters for typographical errors. Correct any errors in the file.

Renaming Groups or Users

When you rename a group or user in the LDAP directory server, the group or user is not renamed in the system. The system cannot identify the object if the primary name of the object has changed, so a new object is created instead. Instead of renaming a group or user, delete it and create one.

Full Name Differs from User Logon Name

The application server synchronizes on cn (common name), which is the **Full Name** field in the LDAP directory server.

To synchronize from the user logon name and log in to the system, the user name attribute must be correctly mapped in LDAPSYNC cron task parameters, VMMSYNC cron task parameters, and 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 using the Database Configuration application (see the *Database Configuration online help*).

For more information about the LDAP active directory and Virtual Member Manager, see the *Installation Guide* for your product offering.

Configuring WebLogic Security for Active Directory

For information about configuring WebLogic Server for LDAP security, see BEA's WebLogic Server documentation and search for *WebLogic Active Directory*:

http://e-docs.bea.com/wls/docs92/

After you complete the procedure to configure WebLogic server to use LDAP security, follow these steps:

1. Complete the tasks in *Configuring the System* on page 48.
2. Restart the application server.
3. Deploy the system and map the security role, **maximouser**, to the users and groups that meet your organization's requirements, or assign the users to the default group, **maximousers**, in the LDAP system.

Configuring WebSphere Security for Active Directory

For information about configuring WebSphere Application Server to use WebSphere security and LDAP security, see IBM's WebSphere Application Server, Version 6.1 Information Center, and search for *WebSphere Active Directory*:

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

After you complete the procedure to configure WebSphere Application Server to use LDAP security, follow these steps:

1. Complete the tasks in *Configuring the System* on page 48, then return to this procedure.
2. Restart the application server.

3 Deploy the system and map the security role, **maximouser**, to the users and groups that meet your organization's requirements, or assign the users to the default group, **maximousers**, in the LDAP system.

For example:

- On the Configuration tab, under the Additional Properties heading, click **Map security roles to users/groups** and select **maximouser**.
  - Click **Look up users** if you want to give individual users access to the system.
  - Click **Look up groups** if you want to give groups and users in groups access to the system.
  - Click **All Authenticated** if you want to give access to all successfully authenticated users.
- Authenticate the users into the system:
  - For individual users, search for **max*** for maximouser or * for all users. Click **>>** to move users from the Available list to the Selected list.
  - For groups, search for **max*** for maximouser groups or * for all users. Click **>>** to move user groups from the Available list to the Selected list.
- Click **OK** when complete and then click **Save** to save the Enterprise Application configuration changes.
- Click **OK** to synchronize changes with nodes.
- In the left pane, go to the **Servers> Application servers** folder and click **maximoserver** and start the server.

Configuring the System

TIP If the application server security was configured through the installer, you do not have to complete these steps.

To configure the system to use application server security, follow these steps:

1 In the system properties file, add the property, `mxe.useAppServerSecurity`. Set the value of the property to 1 and save the file.

2 In `applications\maximo\maximouiweb\webmodule\WEB-INF\web.xml`, uncomment the following lines:

```
<!--
<login-config>
<auth-method>BASIC</auth-method>
<realm-name>MAXIMO Web Application Realm</realm-name>
</login-config>
-->
```

In addition, you can use FORM login. To use FORM login, uncomment this login-config and ensure that the BASIC based login-config (that was previously described) is commented out.

```
<!--
<login-config>
    <auth-method>FORM</auth-method>
    <realm-name>MAXIMO Web Application Realm</realm-name>
    <form-login-config>
      <form-login-page>/webclient/login/
login.jsp?appservauth=true</form-login-page>
      <form-error-page>/webclient/login/loginerror.jsp</form-
error-page>
    </form-login-config>
  </login-config>
-->

<env-entry>
            <description>Indicates whether to use Application
Server security or not</description>
            <env-entry-name>useAppServerSecurity</env-entry-
name>
            <env-entry-type>java.lang.String</env-entry-type>
            <env-entry-value>1</env-entry-value> ====>>>>
</env-entry>
-->
```

3 In this file, set the `useAppServerSecurity` setting to 1.

4 In the same file, uncomment the following lines:

```
<!--
    <security-constraint>
        <web-resource-collection>
            <web-resource-name>MAXIMO UI pages</web-resource-
name>
            <description>pages accessible by authorised users</
description>
            <url-pattern>/*</url-pattern>
            <http-method>GET</http-method>
            <http-method>POST</http-method>
        </web-resource-collection>
        <auth-constraint>
            <description>Roles that have access to MAXIMO UI</
description>
            <role-name>maximouser</role-name>
        </auth-constraint>
        <user-data-constraint>
            <description>data transmission guarantee</
description>
            <transport-guarantee>NONE</transport-guarantee>
        </user-data-constraint>
    </security-constraint>
-->
```

5 Edit the `applications\maximo\mboweb\webmodule\WEB-INF\web.xml` file, and uncomment the following lines:

```
<!--
    <security-constraint>
        <web-resource-collection>
            <web-resource-name>MAXIMO Report Tool</web-resource-
name>
            <description>pages accessible by authorised users</
description>
            <url-pattern>/reporttool/*</url-pattern>
            <http-method>GET</http-method>
            <http-method>POST</http-method>
        </web-resource-collection>
        <auth-constraint>
           <description>Roles that have access to MAXIMO Report
Tool</description>
            <role-name>maximouser</role-name>
        </auth-constraint>
        <user-data-constraint>
            <description>data transmission gaurantee</
description>
            <transport-guarantee>NONE</transport-guarantee>
        </user-data-constraint>
    </security-constraint>

    <login-config>
        <auth-method>BASIC</auth-method>
        <realm-name>MAXIMO Web Application Realm</realm-name>
    </login-config>
-->
```

6 Set the `useAppServerSecurity` setting to 1:

```
<env-entry>
            <description>Indicates whether to use Application
Server security or not</description>
            <env-entry-name>useAppServerSecurity</env-entry-
name>
            <env-entry-type>java.lang.String</env-entry-type>
            <env-entry-value>1</env-entry-value>
     </env-entry>
```

7 Build the EAR file:

   a Open a Command Prompt.

   b Change directory to your <*Maximo root*>\deployment folder, for example:

      `C:\Maximo\deployment`

   c Type **buildmaximoear**.

   d Click **Enter**.

8 Deploy the EAR file in the appropriate application server. For more information about WebSphere Application Servers, see Chapter 12, *System Configuration* on page 223, and search for *Enterprise Application Archive Files*.

9 Synchronize the users and groups from LDAP into the sytem using the cron task. For information about synchronizing data, see *Synchronization* on page 44.

Enabling Auto-creation of LDAP Users

The system has an option to auto-create user records that pass LDAP authentication. This option allows you to create users with basic privileges and later 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`.

Options for User and Group Management when Application Server Security is Enabled

When the property `useappserversecurity` is set to `true`, the following functions are disabled:

- Create users
- Change passwords
- Self registration
- Create security groups
- Associate users and security groups

Directory Owns Group Creation/Group Membership Management

If your implementation does not require that the directory owns group creation and group membership management, set the property, `mxe.LDAPGroupMgmt`, with the default value of `1`. This setting indicates that the directory owns group creation and management.

When you change the value to `0`, the system owns group creation and management, not the directory. This setting enables the following functions:

- Create security groups
- Associate users with security groups

Directory Does Not Own Group Creation/Group Membership Management

If your implementation does not require that the directory owns user creation and management, set the property, `mxe.LDAPUserMgmt`, with the default value of `1`. This setting indicates that the directory owns user creation and management.

When you change the value to `0`, the system owns user creation and management. This setting enables the system to use LDAP for user authentication without having to synchronize user information.

The USERID records created in the directory and the USERID records created in the system must be identical for this setting to function correctly.

	mxe.LDAPUserMgmt = 0
Add/delete security groups	no
Modify security groups	yes
Manage user/group relationships	no
Add/delete users	yes
Modify users (other than password)	yes
User self registration	no
Change password	no

Single Sign On

Single sign on (SSO) is an authentication process that lets a user type one name and password to access multiple applications. When a user authenticates with the server, the single sign on application authenticates the user to access all of the applications to which they have been given rights on the server. This authentication eliminates the need to type multiple passwords when the user switches applications during a particular session.

TIP IBM developed the system with the flexibility to integrate with single sign on systems, but we do not provide software or support for a single sign on system.

The system can participate in a single sign on environment when you enable application server authentication. Both the WebLogic server and WebSphere Application Server support a single sign on environment. Various vendors provide single sign on platforms that are compatible with WebLogic servers and WebSphere Application Servers.

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.

TIP For the system reports to function, configure your report server to support a single sign on environment.

Managing Security Roles

Use the following roles to manage security.

Application Server	Security Role	Description
WebLogic	Supports both global and scoped roles. By default, WebLogic Server uses scoped roles. You can change to global roles using the administration console.	Global - applies to all resources within a security realm (that is the entire server domain). Scoped - Applies to a specific instance of a resource deployed in a security realm.
WebSphere	Supports only scoped roles.	Map the scoped roles to individual non-nested groups. WebSphere Application Server cannot authenticate users in nested groups.

For more information about WebLogic Server, go to the BEA Web site and:

- Search for *securing WebLogic resources*:

  http://e-docs.bea.com/wls/docs81/secwlres/secroles.html#1217798

- Search for *types of WebLogic resources*:

- http://e-docs.bea.com/wls/docs81/secwlres/types.html#1213777

For more information about WebSphere Application Server, go to the IBM Information Center for WebSphere Application Server Network Deployment and search for *security roles*:

- http://publib.boulder.ibm.com/infocenter/ws51help/index.jsp

Authorization

You use the Security Groups application to grant authorizations to security groups. You can define the following authorizations for each security group:

- Access to a Start Center
- Access to one or more Sites
- Read, insert, save, and delete access to applications and access to menu options
- Access to one or more inventory storerooms
- Access to labor records
- Authorization to general ledger components
- Specify purchasing limits and tolerances
- Define conditional data restrictions for objects and attributes

If you implemented LDAP, you create groups and manage groups in the LDAP server (unless you changed the default behavior using the property `useappserversecurity`, as described in *Options for User and Group Management when Application Server Security is Enabled* on page 51.

To create groups and manage groups in the LDAP server, place users in security groups. The combination of these groups represents users' security profiles. Users acquire the authorizations and rights of the security groups to which they belong. Users' security profiles determine their maximum rights and privileges.

Group Access

The application access of a group is linked to its Site access. You can give a group access to:

- All Sites
- Specific Sites
- No Sites

With the Security Groups application, you can set access to storerooms, labor, general ledger components, and approval limits and tolerances. When you create a security group, the following options are available on different tabs:

- Access to all storerooms
- Access to all general ledger segments
- Access to all Sites
- Access to all labor
- Access to all labor in your crew
- Access to all labor in your person group
- Access to all labor you supervise
- Access to your own labor

Application Access

Users have four types of access to an application:

- Read
- Insert
- Save
- Delete

You can grant users specific options within an application. For example, you can grant managers rights to read work order histories, costs and warranties, but not to insert work orders or service requests.

You must minimally grant users read access to applications. 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 these types of column information:

- Application Option Description
- Application Option Name
- Visible

- Also Grants
- Also Revokes
- Prerequisite

The visible setting (**Y** or **N**) indicates whether you can select the option from the Applications authorization tab in the Security Groups application. If an option is not visible, it is granted with another option. These standard system options are not visible:

- Clear
- Bookmark
- Next
- Previous
- Viewhist
- Drilldown

For example, when you select READ, the system grants the invisible CLEAR, BOOKMARK, NEXT, PREVIOUS, VIEWHIST, and DRILLDOWN options. The READ option alone does not provide useful function.

The Also Grants, Also Revokes, and Prerequisite values indicate inter-relationships between options.

Example

- If you select the **INSERT** option for an application, the system Also Grants the SAVE option.
- If you deselect the **SAVE** option, the system Also Revokes the **INSERT**, **DUPLICATE**, and **DELETE** options.

Standard access options are typically associated with **Prerequisite**, **Also Grants**, and **Also Revokes** options:

Standard Prerequisite	Prerequisite
Duplicate	Insert
Delete	Save
Standard Also Grants	**Also Grants**
Insert	Save
Read	Clear, Bookmark, Next, Previous, Viewhist, Drilldown
Standard Also Revokes	**Also Revokes**
Read	All Options
Save	Insert, Duplicate
Insert	Duplicate

These inter-relationships are typically true, but individual applications sometimes vary. To view option access information for a specific application, use a SQL editor to search the SIGOPTION table.

Conditional Security Overview

You can configure the system to meet different user requirements. In the system, there is function that enhances user interface, access to data based on user group, and the evaluation of one or more conditions.

TIP Conditional security is a powerful tool; use it judiciously. Extensive use can degrade system performance.

In the system, you can use data restrictions to limit the data to hide entire records or to make them read-only. You can also create data restrictions at the attribute level and make them hidden, read-only, or required. Because these data restrictions exist at the data level, they apply to any user interface element or application that uses an object or attribute.

The following information can help you determine the best place to create restrictions:

- Data restrictions always win over 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
  - Application Designer application
- Configurations that you create with data restrictions apply everywhere 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 of the fields inherit the restriction. If you configure the same restriction in the Application Designer application, apply the same configuration to each field on each tab.
- Configurations using Application Designer are always for one application. Configurations that use data restrictions can apply either to all applications that use the object/attribute or to one specific application.

TIP A data restriction created on an object does not apply to views of that object. Create a separate restriction for the view.

Conditional Expression Manager

The system has a condition library. Within this library, a user can define conditions, either as expressions or as custom class files that you can use to drive application behavior. You can use conditions within the library for conditional option access, data restrictions, and conditional user interface. The expressions for these conditions use a syntax that is like SQL, but that uses bind variables for reusability.

The Conditional Expression Manager application facilitates the management of these conditions. You can use the Conditional Expression Manager application to define conditions used in data restrictions and elsewhere in the system. Using this application, you can create and maintain a library of conditions.

When the expression evaluates to true, conditions let you configure access to fields, tabs, and other user interface controls within applications. For example, you can set the following types of conditional access:

- Give read-only access to information displayed in a field
- Give read-write access to information displayed 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 a specific application
- Do not display a field or tab in an application to certain users

You can also set access to application options, controls, or data elements. Examples include:

- 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 differ according to group and different conditions
- Show or hide a data attribute globally or for a security group

Defining Conditions

You use the Conditional Expression Manager application to define conditions used in data restrictions. The system uses a syntax like SQL. However, the syntax of the system uses additional variables.

Expression Syntax

The expression syntax of the system uses a colon (:) to define a variable. The system uses this syntax to avoid ambiguity when you create expressions that relate to the current record (business object) or to a specific record. For example, you can have subselect in the expression, such as:

```
exists (select 1 from workorder where wonum=:wonum)
```

The first `wonum` is the wonum attribute on the `workorder` object. At run time, the system replaces the second `wonum` with the value of the `wonum` attribute for the current business object.

Replacement Variables

Syntax	Description	Comments
:yes	true	Logically true, `1` if stored in the database
:no	false	Logically false, `0` if stored in the database
:&date&	Current date	
:&datetime&	Current date/time	
:&user&	Logged in user	For example, if a user is signed in as Smith, `:owner.id=&user&` converts to `:ownerid='SMITH'`

Syntax	Description	Comments
:&personid&	Person ID of the logged in	For example, if a user is logged in as Smith, `:reportby=&personid&` converts to `:reportby='SMITH'`

:&appname&	Application name	For example, in the Work Order Tracking application:
		`':&appname& = WOTRACK'`
		converts to:
		`'WOTRACK = WOTRACK'`
		This variable is useful for setting different behavior for different copies of an application.

:&mboname&	Name of the current business object	For example, in the work order object:
		`'object = :&mboname&'`
		converts to:
		`'object = WORKORDER'`

:&ownername&	Name of the owner business object	For example, in the Work Order Tracking application:
		`:&owner&.jobplan.priority>&owner&.priority`
		converts to:
		`workorder.jobplan.priority>workorder.priority`

Bind Variables

Syntax	Description	Examples
:<relationshipname>.<attrname>	Value of an attribute of a related business object of the current business object	:location.description
:&owner&.<attrname>	Value of an attribute of the owner business object	When you apply a job plan to a work order, the system copies the priority of the job plan to the child work order. The condition can be: `:&owner&.jobplan.priority>:&owner&.priority` In this example, the system copies this information if the job plan has a higher priority than the parent work order.
:&owner&.<relationship_name>.<attrname>	Value of an attribute of the related business object of the owner business object	See the example for :&owner&.<attrname>
:$old_<attrname>	The initial value from the database of the attribute	For example, if you change the value of a field from `1` to `2` to `3` to `4`, the old value is `1`.

Sample Expressions

The system can use the following sample expressions:

```
:wostatus='APPR'
:type='EM'
:ownerid=:&user&
:supervisor!=:&personid&
:asset.assettype = 'IT' and :&personid&=:owner
:reportby=:&personid&
:assetspec.classstructureid = 1221
:po.poline.receivedqty=0
:&owner&.jobplan.priority>:&owner&.priority
:&owner&.po.$old_description like '%Turbin%'
```

Sample Condition Classes

The system can use these condition classes:

- evaluates to true:
  `psdi.common.condition.AlwaysTrue`
- evaluates to false:
  `psdi.common.condition.AlwaysFalse`

These files are in the directory structure, and must be in a folder under the root of the product installation. To use a sample condition class, type the file name in the **Class** field.

Data Restrictions

Using Data Restrictions in the Security Groups application, you can set restrictions on exactly which records a group can access within the larger set specified by the privileges you grant on the Applications tab.

- The Global Data Restrictions action in the Select Action menu lets you use a condition to set restrictions on which records can be accessed in the system.
- The Data Restrictions tab in the Security Groups application lets you use a condition to set restrictions on which records a group can access.

Data restrictions provide you with numerous 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/application with a condition to qualify the data to be returned. In this case, only data meeting the condition is fetched from the database, which differs from data that is fetched from the database, but hidden in a certain condition.
- 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/attributes might change their display as a user scrolls through records.
- You can set collection restrictions to control the collections of Assets, Locations, and configuration items (CIs) that a group can access.

TIP When you grant a user access to an application, the user has access to all of the data elements per the business logic of that application.

Group Data Restrictions

You use the Data Restrictions tab in the Security Groups application to 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 a certain way:

- Qualified data restrictions are ORed together
- 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 in the Select Action menu to set restrictions using a condition that defines which records can be accessed in the system. See *Conditional Expression Manager* on page 56 for additional information about creating and managing expressions.

Conditional User Interface

In the system, you can bind any user interface control to a Signature Option and, thereby, grant or revoke the user interface control to groups of users.

For additional information, see the *Application Developer Guide*.

Encryption

The system uses the data types, Crypto, and CryptoX, to encrypt passwords and other types of confidential information.

Data Type	Data Stored	Algorithm
CryptoX	User passwords	▼ One-way encryption ▼ Stores password in encrypted format (cannot be decrypted or displayed) ▼ Internally, the system uses the encrypted version
Crypto	Information you want to decrypt for display	▼ Two-way encryption ▼ Information can be decrypted and displayed to users

The system uses the Java™ Cryptography Extension (JCE) to perform encryption. This technology can use variables (for example, Provider, Mode, Padding, Key, and Spec) to transform the input data into encrypted data. By default, the system uses the DESede algorithm and its defaults for the other values. Crypto and CryptoX use the DESede encryption algorithm.

Modifying Encryption Settings

As a best practice, modify the encryption settings in the maximo.properties file. You can configure the following encryption data types to be consistent with industry guidelines and government guidelines, or to make your system more secure:

- Key
- Mode
- Padding
- Spec

Encryption Property Name	Settings for JCE and DESede
mxe.security.crypto.key mxe.security.cryptox.key	Length must be a multiple of 24
mxe.security.crypto.mode mxe.security.cryptox.mode	CBC: Cipher Block Chaining Mode
	CFB: Cipher Feedback Mode
	ECB: Electronic Codebook Mode
	OFB: Output Feedback Mode
	PCBC: Propagating Cipher Block Chaining
mxe.security.crypto.padding mxe.security.cryptox.padding	NoPadding
	PKCS5Padding
mxe.security.crypto.spec mxe.security.cryptox.spec	Length must be a multiple of 8

As best practice, modify the encryption settings once. For additional information, see *System Properties* on page 233.

Encrypting Properties

You can encrypt properties to provide additional security. You can encrypt the following properties:

- mxe.db.password
- mxe.system.regpassword
- mxe.report.bo.rptServerLogonPass

  Located in additionalmaximo.properties by default — you must copy it into maximo.properties before you can encrypt it.

The properties reside in the `<Maximo root> \applications\Maximo\properties` folder.

When you encrypt a property, it leaves behind the unencrypted original. For security purposes, store the unencrypted original somewhere outside the system file structure.

To encrypt properties:

1. Use a text editor to modify `maximo.properties`.
2. Open a command shell and go to `<Maximo root> \tools\maximo folder`.

3 Type **encryptproperties** to run the batch file. The old files are renamed with an ***_orig** extension:

- `maximo.properties_orig`
- `ldapsync.xml_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.

Encrypting Additional Properties

You can encrypt additional properties, in addition to the ones listed above.

To encrypt additional properties:

1 Open `encrypt.properties` in a text editor.

2 Add the additional properties that you want to encrypt.

3 Run the encryption procedure described in *Encrypting Properties* on page 62.

The additional encrypted properties must be decrypted wherever they are used in the application. Your development team is responsible for this customization.

Editing Encrypted Files

If you want to edit a file that you already encrypted, follow this process:

- Delete the encrypted `maximo.properties` and `ldapsycn.xml` files.
- Restore the unencrypted originals back into the <*Maximo root*> \applications\Maximo\properties folder.
- Remove the **_orig** extensions from both files.
- Make your changes, then re-encrypt the files.

3 Database Configuration

You use the Database Configuration application to create or modify the objects and attributes, and to customize the database.

The settings specified in the Microsoft® Windows® Control Panel can affect how numbers, currencies, dates, and time are formatted. You can specify the settings for each client computer using the Regional and Language options of the Control Panel.

For additional information, see *Database Configuration online help*.

Data Dictionary

The structure of a relational database is stored in the Data Dictionary of the database.

TIP User error can corrupt the data dictionary. The only way to recover a data dictionary is to restore from a backup.

Database structure lets you interpret data, and to see patterns and trends. When you know how your data is structured, it is easier to retrieve data.

Table Name	List includes
MAXOBJECT	All objects. Links an object to its table or view.
MAXTABLE	All tables.
MAXVIEW	All views.
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 internally generated from this table, but Oracle® and IBM® DB2® use database sequence generators.
MAXSYSINDEXES	All indexes used in the system.

Reserved Words for IBM DB2 Version 8.2

For the most current list, see the IBM Web site and search for *reserved words*.

In addition to reserved words, IBM DB2 uses system-generated names beginning with `SYS_` for implicitly generated schema objects and subobjects. To avoid possible conflict in name resolution, IBM discourages use of this prefix in the names that you explicitly provide to your schema objects and subobjects.

Attributes whose input values must be verified as not being a reserved word are checked against the native database. If the check generates an error, the system concludes that it is a reserved word.

ADD	EXTERNAL	PARAMETER
AFTER	FENCED	PART
ALIAS	FETCH	PARTITION
ALL	FIELDPROC	PATH
ALLOCATE	FILE	PIECESIZE
ALLOW	FINAL	PLAN
ALTER	FOR	POSITION
AND	FOREIGN	PRECISION
ANY	FREE	PREPARE
APPLICATION	FROM	PRIMARY
AS	FULL	PRIQTY
ASSOCIATE	FUNCTION	PRIVILEGES
ASUTIME	GENERAL	PROCEDURE
AUDIT	GENERATED	PROGRAM
AUTHORIZATION	GET	PSID
AUX	GLOBAL	QUERYNO
AUXILIARY	GO	READ
BEFORE	GOTO	READS
BEGIN	GRANT	RECOVERY
BETWEEN	GRAPHIC	REFERENCES
BINARY	GROUP	REFERENCING
BUFFERPOOL	HANDLER	RELEASE
BY	HAVING	RENAME
CACHE	HOLD	REPEAT
CALL	HOUR	RESET
CALLED	HOURS	RESIGNAL

CAPTURE	IDENTITY	RESTART
CARDINALITY	IF	RESTRICT
CASCADED	IMMEDIATE	RESULT
CASE	IN	RESULT_SET_LOCATOR
CAST	INCLUDING	RETURN
CCSID	INCREMENT	RETURNS
CHAR	INDEX	REVOKE
CHARACTER	INDICATOR	RIGHT
CHECK	INHERIT	ROLLBACK
CLOSE	INNER	ROUTINE
CLUSTER	INOUT	ROW
COLLECTION	INSENSITIVE	ROWS
COLLID	INSERT	RRN
COLUMN	INTEGRITY	RUN
COMMENT	INTO	SAVEPOINT
COMMIT	IS	SCHEMA
CONCAT	ISOBID	SCRATCHPAD
CONDITION	ISOLATION	SECOND
CONNECT	ITERATE	SECONDS
CONNECTION	JAR	SECQTY
CONSTRAINT	JAVA	SECURITY
CONTAINS	JOIN	SELECT
CONTINUE	KEY	SENSITIVE
COUNT	LABEL	SET
COUNT_BIG	LANGUAGE	SIGNAL
CREATE	LC_CTYPE	SIMPLE
CROSS	LEAVE	SOME
CURRENT	LEFT	SOURCE
CURRENT_DATE	LIKE	SPECIFIC
CURRENT_LC_CTYPE	LINKTYPE	SQL
CURRENT_PATH	LOCAL	SQLID
CURRENT_SERVER	LOCALE	STANDARD
CURRENT_TIME	LOCATOR	START
CURRENT_TIMESTAMP	LOCATORS	STATIC

CURRENT_TIMEZONE	LOCK	STAY
CURRENT_USER	LOCKMAX	STOGROUP
CURSOR	LOCKSIZE	STORES
CYCLE	LONG	STYLE
DATA	LOOP	SUBPAGES
DATABASE	MAXVALUE	SUBSTRING
DAY	MICROSECOND	SYNONYM
DAYS	MICROSECONDS	SYSFUN
DB2GENERAL	MINUTE	SYSIBM
DB2GENRL	MINUTES	SYSPROC
DB2SQL	MINVALUE	SYSTEM
DBINFO	MODE	TABLE
DECLARE	MODIFIES	TABLESPACE
DEFAULT	MONTH	THEN
DEFAULTS	MONTHS	TO
DEFINITION	NEW	TRANSACTION
DELETE	NEW_TABLE	TRIGGER
DESCRIPTOR	NO	TRIM
DETERMINISTIC	NOCACHE	TYPE
DISALLOW	NOCYCLE	UNDO
DISCONNECT	NODENAME	UNION
DISTINCT	NODENUMBER	UNIQUE
DO	NOMAXVALUE	UNTIL
DOUBLE	NOMINVALUE	UPDATE
DROP	NOORDER	USAGE
DSNHATTR	NOT	USER
DSSIZE	NULL	USING
DYNAMIC	NULLS	VALIDPROC
EACH	NUMPARTS	VALUES
EDITPROC	OBID	VARIABLE
ELSE	OF	VARIANT
ELSEIF	OLD	VCAT
ENCODING	OLD_TABLE	VIEW
END	ON	VOLUMES

END-EXEC	OPEN	WHEN
END-EXEC1	OPTIMIZATION	WHERE
ERASE	OPTIMIZE	WHILE
ESCAPE	OPTION	WITH
EXCEPT	OR	WLM
EXCEPTION	ORDER	WRITE
EXCLUDING	OUT	YEAR
EXECUTE	OUTER	YEARS
EXISTS	OVERRIDING	
EXIT	PACKAGE	

Reserved Words for Oracle Version 9.2

For the latest list, see the Oracle Web site and search for *reserved words*.

In addition to these reserved words, Oracle uses system-generated names beginning with `SYS_` to implicitly generate schema objects and subobjects. To avoid possible conflict in name resolution, Oracle discourages using this prefix in the names that you explicitly provide to your schema objects and subobjects.

Attributes whose input values must be verified as not being a reserved word are checked against the native database. If the check generates an error, the system concludes that it is a reserved word.

ACCESS	IDENTIFIED	RAW
ADD *	IMMEDIATE *	RENAME
ALL *	IN *	RESOURCE
ALTER *	INCREMENT	REVOKE *
AND *	INDEX	ROW
ANY *	INITIAL	ROWID
AS *	INSERT *	ROWNUM
ASC *	INTEGER *	ROWS *
AUDIT	INTERSECT *	SELECT *
BETWEEN *	INTO *	SESSION *
BY *	IS *	SET *
CHAR *	LEVEL *	SHARE
CHECK *	LIKE *	SIZE *
CLUSTER	LOCK	SMALLINT *
COLUMN	LONG	START

COMMENT	MAXEXTENTS	SUCCESSFUL
COMPRESS	MINUS	SYNONYM
CONNECT *	MLSLABEL	SYSDATE
CREATE *	MODE	TABLE *
CURRENT *	MODIFY	THEN *
DATE *	NOAUDIT	TO *
DECIMAL *	NOCOMPRESS	TRIGGER
DEFAULT *	NOT *	UID
DELETE *	NOWAIT	UNION *
DESC *	NULL *	UNIQUE *
DISTINCT *	NUMBER	UPDATE *
DROP *	OF *	USER *
ELSE *	OFFLINE	VALIDATE
EXCLUSIVE	ON *	VALUES *
EXISTS	ONLINE	VARCHAR *
FILE	OPTION *	VARCHAR2
FLOAT *	OR *	VIEW *
FOR *	ORDER *	WHENEVER *
FROM *	PCTFREE	WHERE
GRANT *	PRIOR *	WITH *
GROUP *	PRIVILEGES *	
HAVING	PUBLIC *	

* These words are also ANSI-reserved.

Reserved Words for SQL Server

Microsoft SQL Server 2000 uses reserved keywords for defining, manipulating, and accessing databases. Keywords are part of the grammar that SQL Server uses to parse and understand Transact-SQL statements and batches.

It is syntactically possible to use SQL Server reserved keywords as identifiers and object names in Transact-SQL scripts; however, you must use delimited identifiers.

For a current list, see Microsoft documentation.

ADD	EXCEPT	PERCENT
ALL	EXEC	PLAN
ALTER	EXECUTE	PRECISION
AND	EXISTS	PRIMARY
ANY	EXIT	PRINT
AS	FETCH	PROC
ASC	FILE	PROCEDURE
AUTHORIZATION	FILLFACTOR	PUBLIC
BACKUP	FOR	RAISERROR
BEGIN	FOREIGN	READ
BETWEEN	FREETEXT	READTEXT
BREAK	FREETEXTTABLE	RECONFIGURE
BROWSE	FROM	REFERENCES
BULK	FULL	REPLICATION
BY	FUNCTION	RESTORE
CASCADE	GOTO	RESTRICT
CASE	GRANT	RETURN
CHECK	GROUP	REVOKE
CHECKPOINT	HAVING	RIGHT
CLOSE	HOLDLOCK	ROLLBACK
CLUSTERED	IDENTITY	ROWCOUNT
COALESCE	IDENTITY_INSERT	ROWGUIDCOL
COLLATE	IDENTITYCOL	RULE
COLUMN	IF	SAVE
COMMIT	IN	SCHEMA
COMPUTE	INDEX	SELECT

CONSTRAINT	INNER	SESSION_USER
CONTAINS	INSERT	SET
CONTAINSTABLE	INTERSECT	SETUSER
CONTINUE	INTO	SHUTDOWN
CONVERT	IS	SOME
CREATE	JOIN	STATISTICS
CROSS	KEY	SYSTEM_USER
CURRENT	KILL	TABLE
CURRENT_DATE	LEFT	TEXTSIZE
CURRENT_TIME	LIKE	THEN
CURRENT_TIMESTAMP	LINENO	TO
CURRENT_USER	LOAD	TOP
CURSOR	NATIONAL	TRAN
DATABASE	NOCHECK	TRANSACTION
DBCC	NONCLUSTERED	TRIGGER
DEALLOCATE	NOT	TRUNCATE
DECLARE	NULL	TSEQUAL
DEFAULT	NULLIF	UNION
DELETE	OF	UNIQUE
DENY	OFF	UPDATE
DESC	OFFSETS	UPDATETEXT
DISK	ON	USE
DISTINCT	OPEN	USER
DISTRIBUTED	OPENDATASOURCE	VALUES
DOUBLE VARYING		OPENQUERY
DROP	OPENROWSET	VIEW
DUMMY	OPENXML	WAITFOR
DUMP	OPTION	WHEN
ELSE	OR	WHERE
END	ORDER	WHILE
ERRLVL	OUTER	WITH
ESCAPE	OVER	WRITETEXT

The Database Configuration Menu

The following actions are available from the Select Action menu of the Database Configuration application.

Action	Description
Manage Admin Mode	Lets you configure the database without shutting down the application server.
Delete Object	Marks the object for deletion. Objects are not deleted until you apply configuration changes.
Apply Configuration Changes	Changes are written to a secondary table until you apply them. To apply changes, you shut down the application server and run configdb from the command line. For instructions, see *Configuring the Database* on page 82.
Discard Configuration Changes	Discards changes that are not applied. The content of the Configuration objects, which holds the changes until they are applied, is cleared and reloaded with active Metadata values.
Delete Backup Tables	If applicable, the dialog box presents a list of backed up objects. You can select backup objects to delete, but cannot reconfigure the database until all are deleted.
Update Statistics	Improves database performance by reorganizing indexes.
Refresh Index Tables	Checks the database indexes and reloads into the system index metadata.
Field Length and Format	Used to view or change the **Amount** field format (length and decimal precision) and the **Integer** field and **Smallint** field. Not all cost type fields are amounts. You must change the length of the decimal cost type fields in Attributes.
Services	Enables the creation and management of system service entries that represent Java™ components loaded into the application server on startup. These services support the creation and initialization of other Java business objects.
GL Account Configuration	Specifies the general ledger account code format, including component field lengths and types, and delimiters.
Manage eSig Actions	Enables eSignature on actions within an application. Lists applications and associated actions.
Messages	Enables the creation and management of system messages. Messages can be informational, warning, or error.
Restrict Attributes	Determines whether Maximo® Enterprise Adapter can set values to the attributes.
Add to Bookmarks	Lets you access the current record later from the **List** tab.
Run Reports	Lists the available reports. Select a report title and set parameters, and click **Submit**.

About Objects

An object is a self-contained software entity that consists of both data and functions to manipulate data. Every application is associated with an object.

When you use the Database Configuration application, you interact at the business object level. Internally, the application determines the actions to take on the tables to support the needs of business objects.

A database table stores several objects, and each object has different business rules. For example, the TICKET table defines Incident, Problem, and Ticket business objects.

In addition, a business object can span more than one database table. Views represent objects that span multiple tables. For more information about views, see *Creating Views* on page 80.

With the business object layer, the system tables that you must not modify are hidden from the user interface (however, as an administrator, you can view them). Some tables contain modifiable columns which display the appropriate attributes to correspond to those columns.

Creating or Modifying an Object

An object can be associated with a table or a view; and objects can be persistent or non-persistent. You can create user-defined objects in addition to existing objects (typically for custom applications). You specify the number of columns and their attribute definitions.

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

1. In the Database Configuration application, click **New Object** or select an object from the **List** tab.
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. Type a description for the new object.
4. To use Views, specify a value in the **Extends Object** field.

   After you specify a value, the **View?** check box is selected. If the view joins two tables, you can type the name for the second table in the **Join to Object** field.

5 Complete the appropriate fields:

Field	Description / Procedure
Details Section	
Level*	Specify a value. For definitions of these terms, see *Site and Organization Types* on page 96.
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 non-persistent, nothing is added for attributes, but 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. A database storage partition is the location where a database object is located on a disk. Database storage partitions are called tablespaces in IBM DB2 and Oracle, and called file groups in Microsoft SQL Server. The database administrator must configure the value list DBSTORAGEPARTITION to include a valid list of available tablespaces/file groups. (for example, avoid creating objects in the SYSTEM tablespace.) In IBM DB2, the database or system can manage tablespaces: ▼ If the database manages tablespaces, indexes can be different from the table. ▼ If the system manages tablespaces, indexes must be the same as the table.

Field	Description / Procedure
User Defined?	The object is a regular product object (User Defined = 0) or an administrator created object (User Defined = 1). If the **User Defined?** check box is selected, the value in the **Imported** field indicates how the object was defined: ▼ **Imported** = 0 — the object was defined using the Database Configuration application. ▼ **Imported** = 1 — the object was not defined using Database Configuration. (This field appears only for existing objects.)
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 L_<tablename>.
Text Search Enabled?	Select to enable text search on the object. You can use this function with text search on attributes. (This field appears only for existing objects.)
View Section	
View	Select the **View** check box. You can edit the **View Where** field, **Join to Object** field, **View Select** field, **Automatically Select** field, and **View From** field. For information about views, see *Creating Views* on page 80.
Audit Section	

Field	Description / Procedure
Audit Enabled?	Check to enable electronic audit records. You can: ▼ Edit the default Audit Table name and the E-audit filter field. ▼ Type the E-signature filter field regardless of the Audit setting. For information about these features, see *Electronic Signatures and Audit Records* on page 90.

6 Click **Save Object**.

Object	Status Field
New	To Be Added
Modified	To Be Changed

The object is not modified or the table or view is not created until you configure the database. For additional information, see *Configuring the Database* on page 82.

Saving Changes to the Database

You save changes to the database; however, the changes do not take effect until you configure the database. You can finish modifying all relevant tabs, and then configure the database.

Saving stores your changes in temporary database configuration objects but does not implement them in the database. Before configuring the database, you can close and reopen the Database Configuration application without losing saved changes.

A secondary table stores pending changes, which also appear in the **Status** field. You cannot query on **Status**.

About Attributes

Each database record contains multiple attributes. For example, the ASSET object contains ASSETID, DESCRIPTION, and GLACCOUNT. You can use the Attributes tab to modify existing attributes or add attributes to the database record.

Data Types

Every attribute has an associated data type.

Value	Description
ALN	Alphanumeric characters, mixed case. Maximum length depends on the database: ▼ Oracle = 4000 characters ▼ SQL Server = 8000 characters ▼ IBM DB2 = 32672 characters
AMOUNT	Decimal number, used for currency.
BLOB	Binary Large Object. Stores JPEGs, movies, or PDFs 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 passwords.
CRYPTOX	Encrypted binary (one-way). Encrypts data in the database, but leaves it readable on the screen. Used for password hints.
DATE	Date only, no time.
DATETIME	Date and time.
DECIMAL	Decimal number. A number including an integer and a fraction that consists of a fixed number of digits called the scale.
DURATION	Displays as 1:30 = 1.5 hours
FLOAT	Floating number. Numbers with fractional portions with variable precision.
GL	General Ledger account. An ALN that is used for GL Accounts. Use this type for GL Accounts.
INTEGER	Integer number.

Value	Description
LONGALN	Long Alphanumeric. Oracle LONG is a character type whose max length = (2**31)-1, but in the system, only 32767. ONGALN is used only for non-persistent 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.
YORN	Yes or No (1 or 0 in the database).

Adding Attributes to Objects

You can use an abbreviation of your organization as a prefix to any new object or attribute name, for example, ACME_EXTRATABLE or 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 using the Application Designer function.

See *Database Configuration online help* for information about how to add object views and object tables.

Class Names

The validation class for the **Class** field extends MboValueAdapter or MaxTableDomain to add business rules to the attribute. The extended class:

- Provides field validation
- Gets a list to show a list of values
- Performs an action on the field Using Class Names

Modifying Attributes

Before modifying an attribute, verify whether it was created by the system or by someone at your site (the **User Defined?** box is selected). You cannot delete attributes created by the system.

Attributes created by the system have more restrictions on modifications than do user-defined attributes. 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 and you can make it anything you want.

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

See *Database Configuration online help* for information about how to modify and delete attributes.

Creating Views

Tables can contain many columns and rows. Relevant information includes:

- Only some of the columns
- Only rows that satisfy a certain condition
- Some columns of one table and some columns of a related table

To filter data, you can create a view: a subset of a database that an application can process. The view can contain parts of one or more tables.

A view does not contain data. Instead, the view is a definition that sits in the data dictionary, along 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.

Purpose

Since views are stored as named queries, you can use views to store frequently used, complex queries. Run the queries using the name of the view in a simple query.

View Select is optional; when it is not specified, all columns in both tables are included.

Populating Views

A view is populated depending on the object on which it is based. For example, if a view is based on the object WORKORDER 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, the default value to WORKORDER, the change is not applied to the view. Instead, you must apply this change to the view, if so desired.

About 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 delete indexes and re-create their definitions.

The **Storage Partition** field lets you select a storage partition for an index.

In IBM DB2, the database or system can manage tablespaces:

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

Both the database and the system cannot manage tablespaces; you must chose one or the other.

See *Database Configuration online help* for information about how to add, modify, and delete an index.

Database Relationships

Database relationships are associations between tables, which are like family relationships:

Type of relationship	Description	Analogy
One-to-one	Both tables can have only one record on each side of the relationship. Each primary key value relates to only one (or no) 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.	Spouse + spouse If you are married, you and your spouse each have one spouse.
One-to-many	The primary key table contains only one record that relates to none, one, or many records in the related table.	Parent + you You have only one mother, but she has several children.
Many-to-many	Each record in both tables can relate to any number of records (or no 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.	Siblings If you have several siblings, so do they.

In the Database Configuration application, you can use the Relationships tab to define SQL for joins, and to create relationships between parent and child objects. Use a JOIN to link data from multiple objects; in the system, the parent is the existing object and the child is the object that you are creating.

For example, Parent = MAXUSER, Child = SITE, and Name = DEFSITE means maxuser exists and you want to get the Site for the user's default Site:

```
siteid = :defsite
```

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

See *Database Configuration online help* for information about creating and deleting relationships.

Configuring the Database

It is important to back up your data before configuring the database.

When modifying the database (examples: creating or deleting objects, attributes, or indexes), changes are stored in secondary tables and do not take effect until you configure the database. Backup tables are stored as part of configuration.

Choosing the Configuration Mode

There are two ways to configure the database: command line mode and admin mode. The following table lists the considerations when determining which configuration mode to use:

Configuration Mode	Considerations
Command line mode	▼ Must shut down the application server
	▼ Requires an Information Technology Admin user who has control over the application server
Admin mode	▼ No need to shut down the application server
	▼ Blocks users from the system applications
	▼ Suspends CRON tasks
	▼ Does not allow remote connectivity
	▼ Disables event listeners
	▼ Requires the user to have Admin login security authorizations, which you must assign in the Security Groups application

For Oracle Only

If the schema owner differs from the database user, you cannot configure the database using the values in the maximo.properties file, as in the following example:

```
mxe.db.schemaowner=MAXIMO
mxe.db.user=MAXIMO_USR
```

If this is the case, you must override the values in the maximo.properties file and run ConfigDB with the -u and -p parameters, and specify the actual schema owner and password:

```
configDB -uschemaowner -ppassword
```

Configuring in Command Line Mode

To configure the database in command line mode, complete the following steps:

1. Shut down your application server and wait one minute before configuring the database. (The application server session timestamp updates every 60 seconds.) Otherwise, a message is displayed:

   ```
   MXServer is running. It must be shut down to run ConfigureDB.
   ```

2. Open a command prompt and change the directory to:

   ```
   <Product home directory>\tools\maximo
   for example: C:\maximo\tools\maximo
   ```

3. Complete one of the following steps:

   - To configure the database and restore backup tables, type **configdb**.
   - (Optional) To avoid restoring backup tables, edit the batch file (ConfigDB.bat file in the `\tools\maximo folder`):

     Sometimes the data in the temp tables (`XX+tablename`) must be modified before restoring.

     a. Remove the `-r` parameter.

     b. Save your changes.

     c. Return to the command prompt and type `configdb`.

   To restore backup tables later, see *Restoring Backup Tables* on page 85.

4. If configuration errors occur, work directly in the database to resolve them. Consult the log files for troubleshooting:

   ```
   <root>\tools\maximo\log
   ```

   For example: `C:\maximo\tools\maximo\log`

5. After configuration completes, 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, see *Restoring Backup Tables* on page 85.

Configuring in Admin Mode

To configure the database in admin mode, you must have admin login security authorization, which lets you to log in when the server is in admin mode. You grant this authorization from the Security Groups application:

1. In the Security Groups application, click the Applications tab.
2. In the Applications table window, search for Start Center.
3. In the Options for Start Center table window, select **Identifies whether the user is permitted to login when server is in Admin Mode**.

To configure the database in admin mode, complete the following steps:

1. In the Database Configuration application, select **Select Action > Manage Admin Mode**.
2. In the Turn Admin Mode ON dialog box, if necessary, 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 dialog box, type the appropriate value in the **Reason for Change** field.
5. Click **OK**.

   A dialog box 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 Admin Mode in 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**.

Non-structural Configuration

You can implement partial live database configuration without turning on Admin Mode. For example, if you add a domain/default value, you do not need to turn on Admin Mode to configure the database.

You can use partial live database configuration for:

- Changes that you can apply without disrupting the live MboSets. If there are any pending changes for which you cannot use partial live database configuration, you must perform a full live configuration or a traditional shutdown configuration.
- For changes that are reversible. If a user applies live changes that are later determined to be incorrect, you can perform another live update to remove the changes.

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 instanciated are not revalidated.

Restoring Backup Tables

You need this procedure only if you did not restore your backup tables during a command line configuration. See *Configuring in Command Line Mode* on page 83.

1. Open a command prompt and change directory to:

   ```
   <Product home directory>\tools\maximo
   ```

2. Run **restorefrombackup**.

3. Start the application server.

Changes Unrelated to eAudit

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

Description	Source
Description	MaxObjectCfg.Description
Esig Filter	MaxObjectCfg.EsigFilter

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

Description	Source	Additional Rules to Qualify for Live Update
Description	MaxAttributeCfg.Remarks	
Title	MaxAttributeCfg.Title	
Domain	MaxAttributeCfg.DomainID	
Default Value	MaxAttributeCfg.DefaultValue	

Description	Source	Additional Rules to Qualify for Live Update
Search Type	MaxAttributeCfg.SearchType	For live update, you cannot change to or from Text Search (domainid = SEARCHTYPE, maxvalue = TEXT)
Esig Enabled	MaxAttributeCfg.EsigEnabled	
Can Autonum	MaxAttributeCfg.CanAutonum	
Autokey Name	MaxAttributeCfg.AutokeyName	
Is Positive	MaxAttributeCfg.IsPositive	
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.

Changes Involving eAudit

An audit table is an object that MaxTableCfg.EauditTbname references. For example, if the Person table is audited and its audit table is named A_Person, two rows exist in MaxTable Cfg:

Table Name	Eaudit Tbname
PERSON	A_PERSON
A_PERSON	(null)

For an audit table, if the object you add (MaxObjectGfg.Changed = I) and its base table are eligible, then the audit table is also eligible.

For the base table, changes to the audit-related, object-level parameters in the following table are eligible for live update.

Description	Source	Additional Rules to Qualify for Live Update
Eaudit Filter	MaxObjectCfg.EauditFilter	
Eaudit Enabled	MaxObjectCfg.EauditEnabled	No additional rules. Field validation in the Database Configuration application ensures that when EauditEnabled is turned on, there is a non-null value for EauditTbname.
Eaudit Table Name	MaxTableCfg.EauditTbname	Live update for change of EauditTbanme is supported, if either of the following is true: ▼ 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 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.

Tracking Number of Database Configurations

There is a maxvar called DBVERSION that is only accessible through the system. This maxvar tracks how many times the database has been configured and increments by 1.

Text Search

The text search function uses standard SQL to index, search, and analyze text and documents stored in the system database. Text search lets users search the system databases for the information that they need. Users can find information based on textual, content metadata, or attributes.

- The Text Search dialog box is accessible from the **Search Type** field on the Attribute tab of applications.
- Text search is only allowed for ALN fields, and is designed to search long descriptions or fields that are long data types.
- Full text search is language-specific text search (not string search). Words, not parts of words, are indexed. For example, the result is not *part* if you search for *par*.
- You must flag text search on the object and appropriate attributes. For example, in the ASSET table, the **Description** field and **Long Description** field are text search-enabled.

TIP DB2 is not enabled for text search.

User Queries

Users generate most of the WHERE clauses in system queries from the **List** tab of applications. This powerful feature can produce inefficient SQL. As a system administrator, you can improve the ease of use and convenience for users by setting the appropriate search types for database columns. Using appropriate search types also reduces the load on the database.

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. Alternatively, because a wildcard search might not qualify for an

index search, the search might result in a full table scan. In this case, the search response is slow.

If an object is enabled for text search, the full text searches on its attributes are provided.

Search Type	Entry Required
Exact match	=
Wildcard	% (example: %value%)
Full text	Any combination of the words in the text search

EXACT Search Type

When a user performs a search, the default method to search many Maximo database fields is to use wildcards (SEARCHTYPE = WILDCARD). The WILDCARD search type causes Maximo to construct a condition of the form:

```
column like '%value%'
```

when the user enters `value` in a field on the **List** tab. In wildcard searching, the database engine cannot use indexes. Searching without indexes can result in slower search times, especially on tables with a large number of rows. You can specify a search type of EXACT when word search is not needed. For example, 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. EXACT searches use wildcards only if a user explicitly enters wildcard characters on the **List** tab or in the WHERE clause.

WILDCARD Search Type

Tables with fewer than 2,000 or 3,000 records are typically scanned regardless of indexes. The I/O cost to read the entire table is less than the average I/O cost of the index lookup plus the table lookup. The SEARCHTYPE value has no effect on database behavior when such scans are performed. You can use the default search type of WILDCARD on description fields of tables that have a relatively small number of rows (2,000 or fewer, for example). Wildcard searching provides more flexibility for the users. Tables with relatively few rows have no noticeable degradation in performance.

Full Text Search Type

Most system tables have one or more longer character columns for descriptions, memos, or remarks. You can provide a search type of TEXT, and a corresponding Oracle Text index or SQL Server full text catalog, for columns that have a lot of text. (Full text indexing is not available with the system on IBM DB2®.) Text indexing puts some load on the database because of the constant need for background processing to keep the indexes synchronized. However, text indexing produces efficient word searching of description fields.

Specify a search type of TEXT on description fields for which word searching is required. Use TEXT on description fields of tables with large numbers of rows (tens of thousands, for example). The text search engine takes time to refresh the indexes, so new records might not be found until the text search index refreshes itself. On Oracle, you can modify the procedure `maximo_ts_job_call` 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. (Use SQL Server 2000 Enterprise Manager or SQL Server 2005 Management Studio.)

Stem search is also performed. For example, searching for *service* returns *servicing* and *serviced*.

NONE Search Type

The NONE search type prevents a column from being used in a WHERE clause, and still allows display of the returned values on the **List** tab.

Text Search Syntax

Search phrases that are typed in a text search field can use the following operators:

Operator	Description
-	not included
;	near
~	not (similar to ;)
&	and
| (pipe)	or

You can use these search operators regardless of the database server that the asset management server is using.

- When you use the symbols (such as -, &, and ~) for the search operators, the search is performed with the actual function of the operator.
- If you use the literal description of the operators (such as and, or, near), the search is performed with the descriptions included in brackets, such as {and}.
- If you use search operators that are supported by the database servers, but are not listed in the table above, the operators are included in brackets.
- Search words that are included in double quotation marks are passed as they appear.

Determining What Users Query

Users typically have a well-defined set of columns that they want to query in each application. You can identify these columns in user team interviews during implementation, or later, by examining reports of slow-running SELECT statements. You can then index these columns to improve system performance. Users can create and save their own queries, and can also 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, you can create special-purpose queries (for example, return all PM 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 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 so that, initially, he can see only work orders with SMITH in the Supervisor field.

Restrict Querying in Applications

You can use a combination of methods to control or restrict user access to query features and user ability to query on specific columns:

- 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 - 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 - 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.

Electronic Signatures and Audit Records

Electronic signatures and audit records provide an additional level of security control and auditing capability.

The persons responsible for the content of electronic records can control system access:

Method	Description
Security Groups and Users applications > Security Controls > Login Tracking	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, producing an audit trail

Electronic Signature

Electronic signature provides unique identifiers of users who changed database records or performed actions. The Electronic Signature Authentication dialog box records user names and full user names.

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 enabling electronic signature for a database attribute, the process works as described below:

- When users try to save a change in a field that uses this database attribute, or an implicit save is performed (for example, you click **OK** on the Change Status page in Work Order Tracking), the Electronic Signature Authentication dialog box is displayed.

- Users must complete the appropriate fields in the Authentication dialog box, choose **Select Actions > Manage eSig Actions** and **Save** (or another option).

  All authentication attempts are saved in the LOGINTRACKING object, but authentication must be successful before the system saves the application data.

After enabling electronic signature for an action, the process works as described below:

- When users access the action, the Electronic Signature Authentication dialog box is displayed when users leave that page and dialog box.

- Users must complete appropriate fields in the Authentication dialog box, and authentication must be successful before you can continue with the selected action.

  During authentication, the LOGINTRACKING object records:

  - 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 dialog box)

- Unique transaction identifier

- Key values columns for the record

Electronic Audit Records

After enabling electronic audit records for a database attribute, the process works as described below:

- 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.

- The audit record includes:
  - Copy of the changed data
  - User name of the user who made the change
  - Identifier indicating whether the change was an insert, update, or delete
  - Current date and time of the transaction
  - Rowstamp
  - Unique e-Audit transaction ID
  - 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 (example: the work order number is recorded even when another attribute in the WORKORDER object triggers the electronic audit)

Implementing Electronic Signatures and Audit Records

Using electronic signatures and audit records involves:

- Login tracking
- Electronic signature
- Electronic audit

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).

Enabling Login Tracking

Login tracking lets you specify the number of allowed login attempts and block further attempted logins by a user who exceeds that number. Login tracking also lets you track the number of login attempts and view a user's current login status.

You can use login tracking independently of electronic signature. You must enable login tracking to use electronic signature. For information about how to enable login tracking, see *Security Groups online help* or *User online help*.

Enabling Electronic Signature and Electronic Audit Records on Database Attributes

You can enable electronic signature and electronic audit records independently of one another. However, electronic signatures and electronic audit records 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.

To refine the types of records subject to electronic signature and audit records using the e-Audit/e-Sig Filters, see *E-audit and E-signature Filters* on page 93. See the *Database Configuration online help* for information about enabling electronic signature and electronic audit record.

Enabling Electronic Signature for Accessing Specific Menu Items

You can use the Database Configuration application to enable electronic signature to access specific menu items. See the *Database Configuration online help* for information about enabling electronic signature for menu items.

E-audit and E-signature Filters

To refine the types of information that require these E-audit and E-signature filters, use the E-audit Filter and the E-signature Filter on the Audit section of the Object tab in the Database Configuration application. For more information about how to use the E-audit and E-signature filters, see *Database Configuration online help*.

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 dialog box, which includes a **Reason For Change** field.

- To make the **Reason For Change** field let users type free-form text, no further steps are required.
- 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. See *Adding Values to the Reason For Change Domain* on page 93.

Electronic Signature Authentication

When users perform actions for which electronic signature is enabled, the Electronic Signature Authentication dialog box appears. Users must complete required fields and click **OK**. Authentication must be successful before users can continue.

By default, this dialog box includes the following fields:

Field	Description
User Name (required)	Login ID
full user name (unlabeled)	Read-only, from the DISPLAYNAME attribute in the PERSON object
Password (required)	
Reason for Change*	Type ≤ 50 characters

*Required by default. To customize the screens, use the Application Designer.

Adding Values to the Reason For Change Domain

To add values to the Reason for Change domain:

1. Open the **Domains** application.
2. Open the CHANGEREASON domain.

3 Click **Edit Detail**.

4 In the ALN Domain dialog box, click **New Row**.

5 In the **Value** field and **Description** field, type 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
DELETE	Deleted record

6 Click **OK**.

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

General Ledger Account Configuration

Each General Ledger account code consists of several components (segments). In the Database Configuration application, you define the format of the account code. In the Chart of Accounts application, you specify the valid components to be used. For information about the Chart of Accounts application, see *Chart of Accounts* on page 249.

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 components
- Restrict the number of characters in a component field
- Include a total of ≤ 254 characters/digits

Component Sequence

Account components display in a sequential format, with the farthest left component in the string representing the highest level. For example, the MAXDEMO database includes:

- Component 1 = Cost Center
- Component 2 = Activity

- Component 3 = Resource
- Component 4 = Element

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

component 1	component 2	component 3	component 4
6100	350	SAF	
Cost Center	Activity	Resource	Element

Changing Component Values

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

For 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 3 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 and no configuration is required. However, the GL is now invalid, because the cost center component length was shortened to 3 but has a four digit value (in this example) of 6000.

Required Versus Optional Components

Type of Component	Requires a Value for the Account to be Fully Defined	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) account
  - Has no unknown values (placeholders) in required components
  - Example: 6100-350-SAF is fully defined

- Partially defined (partially specified) account
  - Contains placeholders in some required components
  - Example: 6100-???-SAF (the required Activity component is not specified and therefore contains placeholder characters)

Specifying the General Ledger Account Formats

Each general ledger account code consists of numerous distinct components (also called segments). Use any tab of the Database Configuration application to specify the general ledger account code format.

For information regarding configuring general ledger accounts, see *Database Configuration online help* for more information. For a general discussion of account code formats, see *General Ledger Account Configuration* on page 94.

Site and Organization Types

Site and Organization types apply to an object and describe an object's scope. The following table contains a list of SITEORGTYPES.

Type	Definition, Example
SYSTEM	A System-level object/object. Its security restrictions are applied at the application/object level (in the specific System-level mboset).
SYSTEMORG	A System-level object with Organization as an optional value. These applications are treated like System-level applications, but can ask the Profile for Orgs. `orgid is null or orgid = ...`
SYSTEMSITE	A System-level object/object with Site as an optional value. These objects are treated like System-level applications, but can ask the Profile for a list of Sites. `siteid is null or siteid = ...`

Type	Definition, Example
SYSTEMORGSITE	A System-level object/object, but optionally the record can be linked to an Organization or a Site. Used by the Job Plan application and other future System applications with an optional Orgs andor Sites. These objects are treated like System-level applications, but can ask the Profile for a list of Orgs or Sites. `(siteid in null or siteid = ...) and (orgid is null or orgid = ...)`
SYSTEMAPPFILTER	Used for Users and Groups. Treated like System-level applications, but can ask the Profile for a list of Sites and Organizations in the context of an application so the application can filter data. Filtering is required for Site-level administration of users and groups.
ORG	An Organization-level object/object. The framework applies security for this type. `orgid = ...`
ORGSITE	Treated like Organization-level applications but can ask the Profile for Sites. `(siteid is null or siteid = ...) and orgid = ...`
ORGAPPFILTER	Used for Contracts so the Contract application can filter on its special object rather than using standard security. This and other applications developed as this type are treated as System-level but can ask the Profile for a list of Sites in the context of an application so the application can filter data.
SITE	Site level object. `siteid = ...`
SITEAPPFILTER	Site-level object with application filtering. Reserved for future objects.
ITEMSET	Item Sets. Framework adds the required security restriction. Itemsetid must exist in the user's insert Organization.

Type	Definition, Example
COMPANYSET	Company Set. Framework adds the required security restriction. Companysetid must exist in the user's insert Organization.

Security Issues

Security is applied to all SITEORG types. For certain SITEORG types, you can restrict the result set by appending a condition to the WHERE clause. For example, Site type can be:

```
"siteid=..."
```

Communication Templates

4

Communication templates are available in any application that has the Create Communication action. You use the Communication Templates application to:

- Create and manage generic communication templates that users can use to standardize frequently used e-mail communications (notifications).

  For example, service desk agents can create and send an e-mail message from the Service Requests application, Incidents application, and Problems application, using standardized information from predefined communication templates.

  Recipients can respond, and agents can view the two-way dialog from the Communication Log of the corresponding ticket.

- Create e-mail notifications for use with Workflow and escalation processes.

- Associate file attachments or document folders to templates.

  When communications are sent, any attachments to the template are included in the communication, either in a folder or an individual attachment.

For additional information, see *Communication Templates online help*.

Creating Communication Templates

When you create a communication template, you can specify the following parameters:

- The business object for which the template can be used.

- The applications where the template can be used.

- The address from which the e-mail message is sent.

- The address to which replies are sent.

- The subject line of the message.

- The body of the message.

- One or more recipients of the message. Messages can be sent to roles, persons, person groups, and non-system e-mail addresses.

- Whether each recipient receives the message (To), a carbon copy of the message (CC), or a blind carbon copy of the message (BCC).
- Documents to include as attachments when the message is generated.
- Whether you create an entry in the communication log file.

  If you specify this parameter, whenever a communication template is used to generate a communication, you can view an entry in the communication log.

Communication Templates and Objects

When you create a template, you specify a value in the **Accessible From** field to indicate which applications can use the communication template.

The possible values for the **Accessible From** field are show in the following table.

Value	Description
ALL	Available to users from the **Create Communication** action in all applications, including the escalation and Workflow capabilities
APPS	Available to users from the **Create Communication** action in all applications *except* the Escalations and Workflow Designer applications
ESCALATION	Available for use only with escalations
WORKFLOW	Available for use only with Workflow processes

Default Communication Templates

There are numerous default communication templates. These templates support notification function in the E-mail Listeners Configuration, Escalations, Workflow, Service Requests, and Incidents applications.

The default communication templates are in two categories:

- Templates for the system database (which you can modify to suit your business needs) - You must not delete the default communication templates; they are required for their respective notification function to work.
- Templates for the MAXDEMO database (which you can modify or delete as needed) - You can use the MAXDEMO communication templates in your test environment to gain practical experience with adding and managing templates.

Type of Template	Description
For the database	You can modify, but not delete, these templates; their respective notification functions requires them.
For the MAXDEMO database	You can modify or delete these templates. Use these templates in your test environment to practice adding and managing templates.

For example, the E-mail Listeners application uses several error notification templates. If an error is encountered while staging inbound records, the status of the record is set to ERROR and a notification is sent to the e-mail administrator. In the E-mail Listeners application, you specify the e-mail administrator.

If you did not define an e-mail address for the e-mail administrator, an error is written to Maximo.log, if logging is enabled for this application. The type of error determines the error notification is sent to the e-mail administrator.

For information about error notification templates, see Chapter 6, *E-mail Listeners* on page 109.

Using Templates for Notifications

Workflow

Workflow uses communication templates for notifications. Many individuals encounter a record as it moves through its life cycle. Typically, these individuals want to know about the progress of the record. You can design your Workflow process to generate notifications as required by your business process. Notifications can be made through e-mail or through pager, providing that your paging system supports e-mail.

When Workflow administrators design a Workflow process that includes e-mail notifications, they can:

- Create the notification
- Apply a communication template and modify or complete the notification

For Workflow processes, create templates with role-based recipients. Roles are resolved (example: Purchasing Manager) to an individual, a person group, or e-mail address.

Example

Create a Workflow process for purchase requests:

1. A user submits a request for a laptop. The request enters the Workflow process and waits for approval from an immediate supervisor.

2 The supervisor approves the purchase requisition. The purchase requisition is routed to Finance.

3 When approved, the status is set to approved and the user is notified of the approval.

You can create a template for purchase requisition approvals or rejections, can be sent as the request flows through the Workflow process.

Escalations

You can use escalations to monitor time-sensitive records and key performance indicators (KPIs). When you create escalation records, you can specify that the one or more e-mail notifications are generated when a record reaches the defined escalation point. You can create each notification individually 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 an escalation point defines.

You also can create escalations on the Escalations tab of the Service Level Agreements application.

Example

A service desk agent does not complete assignments within six hours. Configure the system to escalate the assignment to the supervisor (by changing the owner through an action) and notify the supervisor.

Service Desk

Service desk agents can create and send e-mail messages from within the Service Requests application, Incidents application, and Problems application, using predefined communication templates. When you create communication templates for the service desk applications, you ensure that communications with service desk customers contain standardized information.

Users can create and send e-mail messages from other applications in the system.

Notifications using Communication Templates

E-mail notifications include:

- Template ID
- Role or recipient name
- Subject
- Message

If information is sent repeatedly, create a template and attach it as a notification on:

- A node in a Workflow process
- An escalation

Use the Workflow Designer application or Escalations application to create these notifications:

Type of Notification	Description
Free-form	▼ Created without using a communication template ▼ Contain only a subset of the features available in a communication template A template ID is generated, but you cannot reuse it as a template.
Template-based	▼ Created by applying a communication template ▼ Use all the features of a communication template, including attachments The values in the **Role/Recipient** field, **Subject** field, and **Message** field are defaulted from the template. You cannot modify them using the Workflow Designer application or Escalations application.

Using Substitution Variables

When you create a template, you can use substitution variables in the **Subject** field and **Message** field in the e-mail notification. The system resolves the substitution variables that display in the Select Fields dialog box based on the business object that you select in the **Applies To** field.

Example

If the template applies to the object ASSET, the list of variables from which you can select are the column names from the database table/view and the relationship names related to the ASSET object.

When users apply templates and create notifications, the substitution variables from templates are replaced with corresponding values from records that generate notifications.

If the subject line of the template reads:

```
Your Incident ID# is :TICKETID
```

then the ticket number from the incident record replaces TICKETID.

Example: Using Substitution Variables in the Message field

1 In the **Subject** field or **Message** field in the Communications Template application, perform one of the following steps:

- Type a space and a colon before the substitution variable.
- Specify a value in the Select Fields dialog box.

The output is formatted correctly.

2 If more text or other variables follow the variable in the **Subject** field or **Message** field, type a space after the variable.

```
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.

Please review these details and contact us if any information is
incorrect.

Phone: :AFFECTEDPHONE

Problem Description: :DESCRIPTION
```

You can also use dot notation with relationships in substitution variables, such as:rel1. rel2. fieldname (rel1 and rel1 indicate relationship names).

Escalations

5

An escalation is a mechanism to monitor records which can take actions or send notifications when a record reaches a defined escalation point. You use the Escalations application to create, view, modify, and delete escalation records.

You can create an escalation for any business object. Because all applications are associated with business objects, you can create escalations for any application.

For additional information, see *Escalations online help*.

Example Ticket Escalation

The following example illustrates escalations that can be built with the Escalations application.

By default, the network support group owns all tickets related to network issues. If tickets are not resolved within three hours, the service provider:

- Escalates priority to high
- Passes ticket ownership to a supervisor
- Sends an e-mail notification to people within the Organization regarding a danger of service level agreement non-compliance

Escalation Components

An escalation record consists of the following elements:

- 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 or Site — Escalations are 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 that the interval be calendar or date based.

- 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.

Understanding Escalations

The application server contains an escalation engine in the form of CRON task that:

- Drives the escalation process
- Leverages the CRON task function
- Tests all active escalation definitions at a set schedule

To trigger escalations, the engine:

- 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 escalation definitions

Modifying Escalations

All fields on an activated escalation are read-only. To modify an escalation record, you first must deactivate it (see *Escalations online help*).

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.

For instructions on how to delete escalations, see *Escalations online help*.

Enabling Logging for Escalations

You might want to 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, access the Logging application, and for each of the loggers listed in the following table, set the desired log level. If the logger does not exist, follow the instructions in *Logging online help* to create a logger. Verify that the loggers are all active. Once 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.

Escalation Loggers

Logger	Description
crontask	Root logger that generates log statements for all CRON tasks
ESCALATION	Child logger that inherits from Crontask root logger and generates log statements specifically for the escalation engine
sql	Root logger that generates log statements of SQL statements run in the application server
COMMTEMPLATE	Child logger that inherits from Service root logger and generates log statements specifically for communication templates and notifications events

For more information about configuring your logging settings, see Chapter 13, *Logging* on page 229.

Escalation and Service Level Agreement Integration

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.

For additional information, see *Service Level Agreement online help*.

E-mail Listeners

6

You use the E-mail Listeners application to receive and process Service Requests and other types of tickets through e-mail in the form of e-mail messages. Using E-mail Listener function, you can create or update tickets, and indicate whether the status is changed or queried based on specified criteria. The E-mail Listeners application can monitor multiple e-mail accounts to retrieve messages. The application supports embedded and normal message attachments.

The E-mail Listeners application cannot process encrypted or digitally signed e-mail messages. Inform users of this limitation.

E-mail Formats

The E-mail Listeners application can process free form e-mail messages or formatted e-mails. Each type of message has different function.

Free Form

Free form e-mails are in plain text and do not follow any specific structure. The E-mail Listeners application extracts the subject line and body of free form e-mails and uses them to either create a Service Request or 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

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 to perform specific actions, such as changing the status of a business object or querying the business object based on criteria.

To support other business objects, create your own workflow process and associated processing logic. A built-in workflow process that only supports various types of ticket objects is shipped.

Storing 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 in the E-mail Listeners application.

The mail server can control attachment size. Contact your mail server administrator regarding these controls, and to determine the file types allowed on E-mail Listener's mail server. Communicate this information to E-mail Listeners users.

Attached Documents

The E-mail Listeners application processes attachments to an e-mail message and stores the attachments as attached documents. For information about the Attached Documents action, see Chapter 11, *Attached Document Configuration and Administration* on page 163.

Attached Documents Example

For example, Sally tries to print a file and receives an indecipherable error message. She sends a free form e-mail with a screenshot describing the problem to help@support.com, the company site for service desk e-mail requests. The E-mail Listener application retrieves Sally's message 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 Sally's initial e-mail, 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.

Components

The following components work together to provide the E-mail Listeners function:

Component	Purpose
E-mail Listeners application Prerequisite: configure mail servers and e-mail accounts.	The application used to create, modify, and delete e-mail listener configurations.
E-mail listener CRON task	Component that runs continuously on the application server and uses the CRON task infrastructure. This component encapsulates a staging process which processes inbound e-mail through a staging table.

Component	Purpose
Workflow Process	A workflow process that parses e-mail information from the staging table and processes it according to the subject and contents of each message.

How the E-mail Listeners Application Works

These events occur when the mail server receives an incoming e-mail:

1. Mail server polling to retrieve unread e-mails from the designated e-mail account
2. E-mail staging, including:
   - Extracting e-mail content, including attachments
   - Storing content in staging tables and attached documents
   - Launching e-mail processing workflow process
3. Workflow processing, depending on the type of e-mail retrieved:
   - For formatted e-mail:
     - Updated: update an existing ticket and communication log
     - Change Status: Change the status of a ticket or any other business object
     - Query: Query an existing ticket or any other business object using specified criteria
   - For free-form e-mail:
     - New: Create a Service Request and communication log
     - Updated: Create a communication log for the existing Service Request

The E-mail Listeners application sends e-mails in response to all incoming e-mails. The response e-mails serve as e-mails confirming that a desired operation was performed on behalf of the user or that there was an error performing the task. If there is an error, the sender must contact the system administrator to resolve the error.

Polling

The Mail Server polling actions include these events:

- Polls the mail server at a specific frequency. See *E-mail Listener online help* to use the Date Selector to set the schedule.

  The **Schedule** field determines the polling frequency.

- For any e-mail marked as read on the mail server, determine whether to delete the e-mail.
  - Yes: mark the e-mail as deleted on the mail server.
  - No: keep the e-mail on the mail server.

  The e-mail deletion rules that you define depend on how the mail server manages the e-mail account. Configure deletion rules in the E-mail Deletion section of the application.

- For new e-mails on the mail server:

  1. Extract the header and message body for each e-mail.
  2. Extract any attachments.
  3. Move to e-mail staging.
  4. Mark e-mail as **Read** on the mail server.

     The E-mail Listeners application manages read e-mails for e-mail accounts accessed using the POP3 mail protocol. POP3 mail protocol does not support marking an e-mail as read on the mail server.

Queuing

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, meaning that e-mail messages are not always processed as quickly as possible. Therefore, it can be beneficial to switch the listener to a parallel processing mode. To perform this switch, configure a queue and associate the queue with the listener.

A queue is an application server component that can facilitate parallel processing. Java™ Messaging Service (JMS) queues that the underlying Java2 Platform, Enterprise Edition (J2EE) Technology application server provides are used. Once a message is placed in a queue, it can be the message processing component can pick it up in an asynchronous manner. In the J2EE application server, these processing components are called message-driven beans. (MDBs)

You can configure the application server to provide multiple MDBs which process multiple messages in parallel, which increases the speed at which the E-mail Listeners application processes e-mails.

To configure queues, follow the steps outlined in , see "Configuring Queues," on page 116.. Once you set up the queues, you can modify or create an e-mail listener definition to specify queue-based processing, queue name, and queue connection factory name.

Staging

The system stages e-mail messages to save all information required to process the e-mail and initiate Workflow processing.

A staging table stores the attributes of an incoming e-mail message, including recipients (To, CC, BC), sender, subject, and message content. This process creates a record; the Workflow process determines how to process the record.

Managing Staging Records

Using the E-mail Listeners application, you can view e-mail records in staging. This function provides an important benefit to administrators who can review incoming e-mail messages and analyze messages that might not be processed. For unprocessed messages, the application enables administrators to edit the Subject or Message contents and resubmit the e-mail message for reprocessing. If necessary, an administrator can also delete e-mail messages that cannot be processed because they are invalid.

Workflow

You use workflow processes to create steps to guide records for your business process.

The system installs a workflow process called Listener Business Process (LSNRBP), You can modify this process or create a process (see *Workflow Designer online help* for additional information).

The LSNRBP workflow performs functions based on the contents of an e-mail message after the message is stored in the staging table. These functions include:

- Change the status of an existing ticket or other business object that is statusable
- Update an existing ticket
- Create a ticket
- Query any business object in the system and return results

If you modify the LSNBRP workflow process or design a new process, you might have to develop Java code to support the actions that your workflow process needs to perform. We recommend that you use professional services to assist you with such tasks.

Customizing E-mail Listener

The base implementation provides defaults for each listener configuration. To customize e-mail listeners, you can specify the Object Key Delimiter and provide your own Preprocessor implementation.

This information is applicable to the processing of free-form messages only.

Object Key Delimiter

The Object Key Delimiter value identifies the incoming e-mail as an existing ticket.

To change the default (##):

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##).

Object Key Identifier

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).

Preprocessor

The default Preprocessor value is psdi.common.emailstnr.Preprocessor. This Java class:

- Run on the server when the listener recognizes a new e-mail
- Parses the Subject line based on the Object Key Delimiter's value, and adds a value to the Object Key column in the E-mail Listener staging table

The preprocessor indicates whether the e-mail is a new or updated ticket:

- The Preprocessor class extracts the substring bounded by the delimiter characters.
- The preprocessor stores recognized substrings in the Object Key column of the staging table. If no substring is recognized, the column is empty.

Customization Scenario

Other characters can represent the Object Key Delimiter. For example, + is the delimiter, and a user sends 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 delimited used with e-mails in your business environment.

Customizing the Preprocessor

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 these public methods:

- Boolean isNewEmail (String del, String subject)
- String getObjectKey (String del, String subject)

In the custom Java class, implement both methods. Each accepts two parameters:

- Delimiter string
- Subject line string

Method	Description	Base preprocessor implementation	Custom implementation
isNewEmail()	Returns a Boolean value indicating whether the e-mail is for a new or 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, or ▼ 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

Java requires that you declare the custom implementation at the beginning of the file. For example:

```
public class MyPreprocessor implements LSNRPreprocessor
```

To customize the preprocessor implementation:

1. Place the Java class source file into an appropriate Java package where you manage all custom Java code.
2. Build your custom Java code into corresponding class files.
3. Build the Enterprise Archive (EAR).
4. Deploy the newly built EAR into the application server for your code changes to take affect.

Configuring Queues

This procedure provides steps to configure JMS queues for e-mail listeners. Configuring queues is optional and is only recommended when you have high volumes of e-mail messages that must be processed quickly.

WebSphere Application Server Steps

To configure the JMS queues for IBM® WebSphere® Application Server, complete the following steps:

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, type your information login screen, type 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 click **New**.

      Ensure that the application node is ctgNode01.
   b. In the **Server name** field, type `MAXIMOSERVER`, and then click **Next**.
   c. Click **Next** to accept the default server template.
   d. Verify that Generate Unique Ports is selected, and click **Next**.
   e. From Confirm new server, click **Finish**.
   f. Click **Save the changes to the master configuration**.
5. To change the MAXIMOSERVER JVM heap size properties, complete the following steps:
   a. In the navigation pane, click **Servers** and then click **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 512.

f Set the **Maximum Heap** size to 1024.

g Click **OK**.

h Click **Save**.

6 To start the MAXIMOSERVER, select **Servers > Application Servers** and select **MAXIMOSERVER**.

7 Click **Start**.

8 Click **Service Integration > Buses**.

9 In the Buses dialog box, click **New**.

10 To add a new service integration bus, type the following information:

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.

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, and change the threshold value for each child configuration.

11 Click **Next**

12 Click **Finish**.

13 Click **Save** to extend the JMS bus setup to the cluster configuration.

14 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.

Adding a Server to the JMS Bus

1. From the WebSphere administrative console, click **Service Integration > Buses**.
2. In the Buses dialog box, click **lsnrjmsbus to open the Buses**.
3. Under Topology in the lsnrjmsbus dialog box, click **Bus members**.
4. In the Bus members dialog box, click **Add**.
5. In the Add a new bus member dialog box, select the server name **ctgNode01:MAXIMOSERVER** to add to the bus.
6. Click **Next**.
7. Select **File store**, and then click **Next**.
8. In the Provide the message store properties panel, click **Next**.
9. Click **Finish**.
10. Click **Save**.
11. Select **lsnrjmsbus**.
12. In the **High message threshold** field, change the value to a minimum value of 500,000 messages.
13. Click **Apply**.
14. Select **Synchronize changes with Nodes**.
15. Click **Save**.

Creating the JMS Bus Destination for the Listener Inbound Queue

To add a logical address for the listener inbound bus destination queue, `lsnrqin`, within the JMS bus, complete the following steps:

1. From the WebSphere administrative console, click **Service Integration > Buses**.
2. In the Buses dialog box, click **lsnrjmsbus**.
3. In Destination resources in the lsnrjmsbus dialog box, click **Destinations**.

4 In the Destinations dialog box, click **New**.

5 In the Create new destination dialog box, verify that queue is the destination type, and click **Next**.

6 In the Create new queue dialog box, type `lsnrqin` in the **Identifier** field and Listener Queue Inbound in the **Description** field.

7 Click **Next**.

8 In the Create a new queue for point-to-point messaging dialog box, select **Node=ctgNode01:Server=MAXIMOSERVER** as the bus member to store and process messages for the lsnrqin bus destination queue.

9 Click **Next**.

10 In the Confirm queue creation dialog box, click **Finish** to complete the creation of the lsnrqin bus destination queue.

11 Go to the path **Buses > lsnrjmsbus > Destinations**, and click **lsnrqin**.

12 In the configuration dialog box, make the following changes

   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.

13 Click **Apply**.

14 Click **Save**.

15 Select **Synchronize changes with Nodes** and click **Save**.

Creating the JMS Connection Factory

You add a connection factory to create connections to the associated JMS provider of point-to-point messaging queues.

To create a JMS connection factory, complete these steps:

1 From the WebSphere administrative console, click **Resources > JMS > Connection Factory**.

2 From the Scope list, select **Cell=ctgCell01**.

3 Click **New**.

4 Verify that the Default Messaging Provider is selected and click **OK**.

5 Type the following information:

Name field: `lsnrconnfact`

JNDI name field: `jms/maximo/lsnr/lsnrcf`

Bus name field: `lsnrjmsbus`

6 Click **OK**.

7 Click **Save**.

8 Select **Synchronize changes with Nodes**.

9 Click **Save**.

Creating the Listener Inbound JMS Queue

You must create a JMS queue, `lsnrqueue`, as the destination for listener inbound point-to-point messages.

To create a listener inbound JMS queue, complete the following steps:

1 From the WebSphere administrative console, click **Resources > JMS > Queues**.

2 From the Scope list, select **Cell=ctgCell01**.

3 Click **New**.

4 Verify that the Default Messaging Provider is selected and click **OK**.

5 Type the following information:

Name: `lsnrqueue`

JNDI name field: `jms/maximo/int/lsnr/qin`

Bus name field: `lsnrjmsbus`

Queue name field: `lsnrqin`

6 Click **OK**.

7 Click **Save**.

8 Select **Synchronize changes with Nodes** and click **Save**.

Creating JMS Activation for the Listener Inbound Queue

You must activate the listener inbound queue, `lsnrqueue`, before the queue can receive messages. Complete the following steps to activate the `lsnrqueue` queue:

1. From the WebSphere administrative console, click **Resources > JMS > Activation Specifications**.
2. From the Scope list, select **Cell=ctgCell01**.
3. Click **New** to complete the General Properties section for the new JMS activation specification.
4. Click **OK**.
5. Type 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`
6. Click **OK**.
7. Click **Save**.
8. Select **Synchronize changes with Nodes** and click **Save**.
9. Stop all IBM-related processes and daemons.
10. Restart these processes for the update to take effect.

WebLogic Server Steps

To configure the JMS queues for the BEA® WebLogic® Server, complete the following steps:

1. Start the WebLogic Server.
2. In the Domain Structure, expand Services and click **Persistent Stores**.
3. Click **New** and select **Create FileStore** to create a File store.
4. In the **Name** field, type the value:

   lsnrstore
5. For the target, accept the default value of AdminServer.
6. In the **Directory** field, specify a folder on the application server computer from which the application server can manage in the store.

The WebLogic Server should be able to perform read and write operations into the store in this folder. For example, a value for a Windows environment can be `c:\tmp`.

7 Click **Finish**.

The system displays the following confirmation message:

File store created successfully.

8 In the Domain Structure, expand **Services** and then expand the **Messaging** entry.

9 In Messaging, select **JMS Servers**, and click **New** to create a JMS server.

10 Name the server lsnrserver and select **lsnrstore** as the persistent store.

11 In the **Target** field, select **AdminServer**.

12 Click **Finish**.

The system displays a confirmation message:

JMS Server created successfully.

13 In the Domain Structure, click **JMS Modules**.

14 Click **New** to create a JMS module.

15 In the **Name** field, type:

lsnrjmsmodule

Leave the **Descriptor File Name** field and **In Domain** field blank. The application server assigns default values.

16 Click **Next**.

17 Select the **Admin Server** check box.

18 Click **Next**.

19 Select the **Would you like to add resource to this JMS system module?** check box.

The system displays the following confirmation message:

The JMS module was created successfully.

20 On the Configurations tab on the Settings for lsnrjmsmodule page, click **New** in the Summary of Resources table.

21 Select **Connection Factory**.

22 Click **Next**.

23 In the **Name** field, type:

lsnrconnfact

24 In the **JNDI Name** field, type:

jms/maximo/lsnr/lsnrcf

25 Click **Next**.

26 Verify that the **Targets** field has the following value:

AdminServer as selected

27 Click **Finish**.

The application server displays a confirmation message:

Connection factory created successfully.

28 On the Configurations tab on the Settings for lsnrjmsmodule page, click **New** in the Summary of Resources table.

29 Select **Queue**.

30 Click **Next**.

31 In the **Name** field, type:

lsnrqueue

32 In the **JNDI Name** field, type:

jms/maximo/int/lsnr/qin

33 In the **Template** field, accept the default value of **None**.

34 Click **Next**.

35 In Targets, select **lsnrserver**.

36 Click **Finish**.

The application server displays a confirmation message:

The JMS Queue was created successfully.

37 On the Configurations tab on the Settings for lsnrjmsmodule page, click the **lsnrconnfact resource**.

38 On the Configurations tab on the Settings for lsnrconnfact page, click the **Transactions** tab.

39 Select the **XA Connection Factory Enabled** option.

40 Click **Save**.

41 In the Change Center, click **Activate Changes**.

Modifying Deployment Descriptors

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. To configure the Message Driven Bean, complete the WebSphere Application Server steps or WebLogic Server steps to remove comment lines from specific sections within the deployment descriptor files of the system.

WebSphere Application Server Steps

1. In your installation folder, locate the deployment descriptor file called `ejb-jar.xml` under the following 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:

   ```
   <!-- 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. Remove the comment lines (<!-- and -->).

   c. Locate the following section:

   ```
   <!-- 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>-->
   ```

   d. Remove the comment lines (<!-- and -->).

3. Save the changes that you made to the file.

4. Locate the file called `ibm-ejb-jar-bnd.xml` under the following file path:

   ```
   /applications/maximo/mboejb/ejbmodule/META-INF folder
   ```

5. Open the file in a text editor and make the following changes:

   a. Locate the following section:

```
<!-- 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>-->
```

b Remove the comment lines (<!-- and -->).

c In the section in which you removed the comment lines, change the value of the <connection-factory-jndi-name> tab to:

jms/maximo/lsnr/lsnrcf

6 Save the changes that you made to the file.

7 Rebuild and redeploy the EAR, as described in Chapter 12, *System Configuration* on page 223 and search for *Enterprise Application Archive Files*.

WebLogic Server Steps

1 In your installation folder, locate the file called `ejb-jar.xml` under the following 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:

```
<!--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 Remove the comment lines (<!-- and -->).

c Locate the following section:

```
<!-- 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>
-->
```

d Remove the comment lines (<!-- and -->).

3 Save the changes that you made to the file.

4 Locate the file called weblogic-ejb-jar-bnd.xmi under the following file path:

```
applications/maximo/mboejb/ejbmodule/META-INF
```

5 Open the file in a text editor and make the following changes:

a 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>
-->
```

b Remove the comment lines (<!-- and -->).

6 Save the changes that you made to the file.

7 Rebuild and redeploy the EAR (see Chapter 12, *System Configuration* on page 205 and search for *Enterprise Application Archive Files*.

Configuring a Chosen E-mail Listener to Use a Queue

Configuring a chosen e-mail listener is the final task in the configuration of JMS queues. To configure an e-mail listener to use a JMS queue, follow these steps:

1 Access the E-mail Listeners application.

2 Select the particular e-mail listener definition that you want to configure to use JMS queue.

Verify that you deactivated the definition.

3 In the Listener tab, perform the following steps:

- Check the **Queue-based Processing?** field
- Clear any default values that are placed in the **Queue Connection Factory** field and **Processing Queue** field.
- In the **Queue Connection Factory** field, type `jms/maximo/lsnr/lsnrcf`.
- In the **Processing Queue** field, type `jms/maximo/lsnr/qin`.

4 Click **Save**.

5 Activate the e-mail listener definition (see the *E-mail Listeners online help* for information about now to activate an e-mail listener definition).

Configuring Security

Overview

Since the E-mail Listeners application enables users to create, update, query, and change the status of tickets, it is critical to ensure that only authorized users can run these functions using e-mails. For the sender of an e-mail, the application checks security authorizations against the system's security configuration. This check establishes the ability of the sender to run each specific function.

A basic requirement to process e-mails is the Person record in the system. Further processing of an e-mail occurs only after the successful location of the Person record using the e-mail address of the sender.

- If a Person record exists and is active, the system locates the corresponding User record.
- If a Person record does not exist or is inactive, the system fails to process and an error response mail is sent back to the user as well as the administrator.
- If the system locates the User record, the system applies the authorizations associated with that user record when the E-mail Listeners application performs security checks on incoming e-mail messages.
- If the system cannot locate the User record, the system uses the `Run As user` of the e-mail listener's underlying CRON task instance.

Security Scenarios

This table describes the two security authorization scenarios that are supported when a user ID exists in the system

Scenario	Support
User exists and has authorization to perform the operation specified in the e-mail	When the listener performs the security check based on the sender of an incoming e-mail. Once the system locates the user record for this sender, the listener constructs a security profile of the user to determine authorizations. If the user has authorizations to perform add, update, and change status operations, the e-mail is processed accordingly.
User exists and does not have authorization to perform the operation specified in the e-mail	In this scenario, the user can update or query only records that the user created.

Assigning Security Authorizations

To assign appropriate authorizations to the users who send formatted e-mails to the system, you must configure security settings using the Security Groups application. For more information about security, see Chapter 2, *Security* on page 11.

Additional Tasks

See the *E-mail Listeners online help* for a description of the configuration process and instructions to activate a listener.

Logging

You can configure the appropriate log levels to get detailed processing and error information from the system's log file regarding the E-mail Listeners application processing of e-mails. See the *Logging online help* for more information about how to use logging with the E-mail Listeners application.

Bounced E-mail

Outbound e-mail that cannot be delivered is called bounced e-mail. Large volumes of bounced e-mail create excess network traffic and affect the ability of the application to process legitimate tickets.

The mail server generates and returns delivery failed messages to the e-mail listener account specified in the **Send From** field, and e-mail listener treats these messages as service requests.

The recommended approach to handle bounced e-mails is:

1. Create a dedicated e-mail account for bounced e-mail notifications to preserve the integrity of the primary e-mail listener account.
2. Base any outbound e-mail notification on communication templates in which the **Send From** field in the template specifies the dedicated bounced e-mail account.

   An e-mail is generated and sent to that address.

Communication Templates that E-mail Listeners Uses

The E-mail Listeners application uses numerous communication templates to generate notifications to end users and administrators. E-mail listener generates the following types of notifications:

- Confirmation - The E-mail Listeners application successfully processed the incoming e-mail message.
- Validation/processing error - The E-mail Listeners application cannot process the incoming e-mail message due to:

- Incorrect formatting
- Incomplete or invalid information in the incoming e-mail message

▼ System error - The E-mail Listeners application encountered a system error while polling or processing inbound e-mail messages.

Confirmation Notifications

Typically, the recipient of confirmation notifications, validation error notifications, and processing notifications is the end user who sent an e-mail for processing. The recipient of system error notifications is typically the system administrator.

The following table lists the communication templates used to generate confirmation notifications after the processing of incoming e-mail.

Template Used for Notification	Description	Recipients
LSNRBPCBSR	Confirmation of creation of a Service Request based on the incoming free-form e-mail	Sender of original e-mail
LSNRBPCHST	Confirmation of status change of records based on the incoming e-mail	Sender of original e-mail
LSNRBPQRY	Response e-mail containing details of results of query from an incoming e-mail	Sender of original query e-mail
LSNRBPUBSR	Confirmation of update of a Service Request based on incoming free-form e-mail	Sender of original e-mail
LSNRBPUOBJ	Confirmation of update of specified object based on the incoming e-mail	Sender of original e-mail
LSNRBPCOBJ	Confirmation of creation of specified object based on the incoming e-mail	Sender of original e-mail

Validation/Processing Error Notifications

The following table lists the communication templates used to generate validation/processing error notifications.

Template Used for Notification	Description	Recipients
LSNRAUTH	Validation error notification when the sender does not have authorization to perform the operation specified in the incoming e-mail	Sender of original e-mail and system administrator
LSNRBPAUTO	Processing error notification when the e-mail listener cannot create an auto key for an attribute that was declared as auto key in the incoming e-mail	Sender of original e-mail
LSNRBPDATE	Validation error when the sender did not specify a properly formatted date or date time value for an attribute that was specified in the incoming e-mail	Sender of original e-mail and system administrator

Template Used for Notification	Description	Recipients
LSNRBPINV	Processing error when the sender specified an update operation for an existing record in the incoming e-mail, but the record does not exist	Sender of original e-mail
LSNRBPUACN	Processing error when the sender specified an invalid action in the incoming e-mail	Sender of original e-mail
LSNRBPUNOB	Processing error when the sender specified an object in the incoming e-mail that the E-mail Listeners application does not support	Sender of original e-mail
LSNRFNKEY	Processing error when sender did not provide all the primary keys for a record that the sender specified in the incoming e-mail	Sender of original e-mail
LSNRFNREQ	Processing error when the sender did not provide all the required attributes for a record that the sender specified in the incoming e-mail	Sender of original e-mail
LSNRINVM	Validation error when the incoming e-mail contains an empty subject line or a subject line that exceeds the allowed length (this error can occur only for free-form e-mails)	Sender of original e-mail
LSNRNOPER	Validation error when the sender of the incoming e-mail does not have a corresponding Person record	Sender of original e-mail and system administrator
LSNRSECAPP	Validation error when the sender of the incoming e-mail does not have the requisite authorizations to perform the operation on the object specified in the e-mail message	Sender of original e-mail and system administrator
LSNRWFMT	Processing error when the incoming e-mail has invalid formatted content	Sender of original e-mail and system administrator

System Error Notifications

The following table lists the communication templates used to generate system error notifications.

Template	Description	Recipients
LSNRBPEX	E-mail listener encountered an error processing the incoming e-mail This error message is generic	Depending on the error, sender of the original e-mail or system administrator
LSNRBPQERR	E-mail listener encountered an error placing information about an incoming e-mail into JMS queue	System administrator

Template	Description	Recipients
LSNRCFGERR	E-mail listener encountered an error when using the listener configuration to connect to an e-mail account	System administrator
LSNRCONNF	E-mail listener encountered an error connecting to the mail server to access the configured e-mail account	System administrator
LSNRERROR	E-mail listener encountered an error processing the incoming e-mail	Depending on the error, sender of the original e-mail or system administrator
LSNRINBF	E-mail listener encountered an error staging the contents of the incoming e-mail in its internal staging table	System administrator
LSNRJMSCF	E-mail listener encountered an error attempting to connect to the configured JMS queue The queue connection factory information for the queue is incorrect	System administrator
LSNRJMSF	E-mail listener encountered an error connecting to the configured JMS queue Error occurs then the queue configuration is correct, however the E-mail Listener cannot establish a connection to the queue	System administrator
LSNRJMSQ	E-mail listener encountered an error connecting to the configured JMS queue Error occurs when the queue information is incorrect	System administrator
LSNRJMMSSYN	E-mail listener encountered an error connecting to the configured JMS queue Further processing of e-mails performed without use of the JMS queue	System administrator
LSNRMAILER	E-mail listener encountered an error when retrieving e-mail from the configured e-mail account	System administrator

Composing Formatted E-mails

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 which keywords you can use when using attribute-value pairs or XML, respectively.

Keyword	Format Type	Required?	Purpose
#MAXIMO_EMAIL_BEGIN	Attribute-value pairs	Yes	Keyword marks the beginning of formatted e-mail content in an e-mail message
#MAXIMO_EMAIL_END	Attribute-value pairs	Yes	Keyword marks the end of formatted e-mail content in an e-mail
LSNRAPPLIESTO	Both	Yes	A value must be provided for this keyword and the value represents a business object upon which an operation is to be performed
LSNRACTION	Both	Yes	A value must be provided. The value can be 'CREATE', 'UPDATE', 'CHANGESTATUS' or 'QUERY'. The value specifies the desired operation to be performed on the business object.
LSNRWHERECONDITION	Both	No	Keyword is used only in e-mails that query a business object. The value represents a valid SQL where condition to be applied on the business object.
LSNRRESULTCOLUMNS	Both	No	Keyword is used only in e-mails that query a business object. The value represents one or more columns from the business object whose values are returned in the response e-mail to the query.
&AUTOKEY&	Both	No	Keyword is used as a value for an attribute or an XML tag specifically for e-mail messages that are intended to create a business object. If this keyword is used, the particular attribute is auto-keyed using standard system function
&SYSDATE&	Both	No	Keyword is used as a value for an attribute or an XML tag specifically for e-mail messages that are intended to create a business object. 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 Composing Formatted E-mails

You must follow general rules when you compose e-mails using attribute-value pairs or XML formatting.

Rules for Attribute-value Pairs Formatting

- To ensure consistent processing of attribute-value e-mails, every e-mail message must contain the `#MAXIMO_EMAIL_BEGIN` keyword and `#MAXIMO_EMAIL_END` keyword. If these keywords are not included in the e-mail message, the system does process the e-mail and sends an error response e-mail to the sender.

- If an e-mail message contains the `# MAXIMO_EMAIL_BEGIN` keyword and `# MAXIMO_EMAIL_END` keyword, the message is treated as attribute-value formatted. If only the `#MAXIMO_EMAIL_BEGIN` keyword or `# MAXIMO_EMAIL_END` keyword is included in the e-mail, the system treats the e-mail as free-form text. Similarly, if the `#MAXIMO_EMAIL_BEGIN` keyword and the `# MAXIMO_EMAIL_END` keyword are not included in the e-mail, the system treats the e-mail as free-form text.

- The syntax for the attribute-value pairs is:

  ```
  Field Title#Attribute Name=Value
  ```

- 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. No other text or character can appear with the semi colon.

- 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 database column name).

- For formatting convenience, an incoming e-mail message can contain both Field Title and Attribute Name separated by #, only Field Title, or only Attribute Name. If only Field Title is provided, the listener attempts to map the title to the appropriate attribute name before processing the message. If the listener cannot map the title or resolve the attribute name, the system does not process the message, and sends an error response mail to the sender.

- All of the attribute-value pairs specified in the e-mail message must occur together. The e-mail listener ignores any other text arbitrarily typed in the e-mail message. To ensure proper processing, demarcate the attribute-value pairs in the e-mail message with the `#MAXIMO_EMAIL_BEGIN` keyword and `#MAXIMO_EMAIL_END` keyword. You must place these keywords on a separate line and end with a new line.

- You can place inline attachments before or after the `#MAXIMO_EMAIL_BEGIN` keyword and `#MAXIMO_EMAIL_END` keyword.

Rules for XML Formatting

- The XML content of an e-mail message must contain a root element of the form `<MAXIMOEMAILCONTENT></MAXIMOEMAILCONTENT>`. All other XML tags and values must be placed within this root element and are treated as children of this root element. This root element is the equivalent of the `#MAXIMO_EMAIL_BEGIN` and `#MAXIMO_EMAIL_END` in Attribute-value pairs.

- The syntax for XML is:

  ```
  ATTRIBUTE NAME attribute='Field Title'></ATTRIBUTE NAME
  ```

- For formatting convenience, an incoming e-mail message can contain only the Attribute Name tag without 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 INBOUNDCOMM table is set to ERROR status and the system sends an error notification to both the end user and the administrator.

- When processing an XML-formatted e-mail message, the listener attempts to resolve the Attribute Name only. If the listener cannot resolve the Attribute Name, the system does not process the message fail and sends an error response e-mail to the sender.

- Because 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 e-mail message body (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. Such reserved characters must be escaped using either standard escape sequences or CDATA constructs.

- The E-mail Listeners application ignores any text outside the XML content, as demarcated by `<MAXIMOEMAILCONTENT></ MAXIMOEMAILCONTENT >`.

- If the keywords &AUTOKEY& or &SYSDATE& must be used in an XML formatted e-mail message, these keywords must be escaped using standard XML CDATA constructs. For example:

  ```
  <TICKETID><![CDATA[&AUTOKEY&]]></TICKETID>
  ```

Examples of Formatted E-mails

The following are examples of formatted e-mails that you can use as reference or as templates to create formatted e-mails to use with the E-mail Listeners application:

Examples of QUERY E-mails

These e-mails examples demonstrate how to use the QUERY function:

Attribute-value pairs query: Query a single record using criteria (`LSNRWHERECONDITION` keyword)

```
#MAXIMO_EMAIL_BEGIN
LSNRACTION=QUERY
;
LSNRAPPLIESTO=SR
;
LSNRRESULTCOLUMNS=TICKETID,DESCRIPTION,REPORTEDBY,COMMODITYGROUP
;
LSNRWHERECONDITION=TICKETID='1001' AND SITIED ='BEDFORD'
;
#MAXIMO_EMAIL_END
```

Attribute-value pairs query: Query a single record without using criteria (`LSNRWHERECONDITION` keyword)

```
#MAXIMO_EMAIL_BEGIN

LSNRACTION=QUERY
;
LSNRAPPLIESTO=SR
;
TICKETID=1002
;
LSNRRESULTCOLUMNS=TICKETID,DESCRIPTION,REPORTEDBY,COMMODITYGROUP
;
#MAXIMO_EMAIL_END
```

Attribute-value pairs query: Query multiple records and returning selected columns

```
#MAXIMO_EMAIL_BEGIN
LSNRACTION=QUERY
;
LSNRAPPLIESTO=SR
;
LSNRRESULTCOLUMNS=TICKETID,DESCRIPTION,REPORTEDBY,INTERNALPRIORITY
,REPORTDATE
;
LSNRWHERECONDITION=STATUS ='CLOSED'
;
#MAXIMO_EMAIL_END
```

Attribute-value pairs query: Query multiple records and returning all columns

```
#MAXIMO_EMAIL_BEGIN
LSNRACTION=QUERY
;
LSNRAPPLIESTO=INCIDENT
;
LSNRRESULTCOLUMNS=*
;
LSNRWHERECONDITION=REPORTEDBY='LIBERI'
;
#MAXIMO_EMAIL_END
```

XML query: Query a single record using criteria (LSNRWHERECONDITION tag)

```
<MAXIMOEMAILCONTENT>
<LSNRACTION>QUERY</LSNRACTION>
<LSNRAPPLIESTO>PROBLEM</LSNRAPPLIESTO>
<LSNRRESULTCOLUMNS>ticketid,description,reportedby,affectedperson,
commoditygroup</LSNRRESULTCOLUMNS>
<LSNRWHERECONDITION>ticketid in('1001')</LSNRWHERECONDITION>
</MAXIMOEMAILCONTENT>
```

XML query: Query without any criteria

```
<MAXIMOEMAILCONTENT>
<LSNRACTION>QUERY</LSNRACTION>
<LSNRAPPLIESTO>PROBLEM</LSNRAPPLIESTO>
<TICKETID>1003</TICKETID>
<LSNRRESULTCOLUMNS>ticketid,description,reportedby,affectedperson,
commoditygroup</LSNRRESULTCOLUMNS>
</MAXIMOEMAILCONTENT>
```

XML query: Query multiple records and returning selected columns

```
<MAXIMOEMAILCONTENT>
<LSNRACTION>QUERY</LSNRACTION>
<LSNRAPPLIESTO>PROBLEM</LSNRAPPLIESTO>
<LSNRRESULTCOLUMNS>ticketid,description,reportedby,affectedperson,
commoditygroup</LSNRRESULTCOLUMNS>
<LSNRWHERECONDITION>AFFECTEDPERSON =’RAMSDALE’ AND STATUS =
‘QUEUED’ </LSNRWHERECONDITION>
</MAXIMOEMAILCONTENT>
```

XML query: Query multiple records and returning all columns

```
<MAXIMOEMAILCONTENT>
<LSNRACTION>QUERY</LSNRACTION>
<LSNRAPPLIESTO>INCIDENT</LSNRAPPLIESTO>
<LSNRRESULTCOLUMNS>*</LSNRRESULTCOLUMNS>
<LSNRWHERECONDITION>AFFECTEDPERSON =’SMITH’ AND STATUS = ‘QUEUED’
</LSNRWHERECONDITION>
</MAXIMOEMAILCONTENT>
```

Examples of CREATE and UPDATE E-mails

These e-mails are examples of how to use the CREATE and UPDATE functions:

Attribute-value pairs create: 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
```

XML create: Create a service request

Body:

```
<MAXIMOEMAILCONTENT>
<LSNRACTION>CREATE</LSNRACTION>
<LSNRAPPLIESTO>SR</LSNRAPPLIESTO>
<TICKETID><![CDATA[&AUTOKEY&]]></TICKETID>
<CLASS>SR</CLASS>
<DESCRIPTION>My XML SR creation e-mail Test</DESCRIPTION>
</MAXIMOEMAILCONTENT>
```

Attribute-value pairs update: Update specific attributes of an existing service request

Body:

```
#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 update: 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 CHANGE STATUS E-mails

These e-mails are examples of how to use the CHANGE STATUS function:

Attribute-value change status: Change status of an existing service request

Body:

```
#MAXIMO_EMAIL_BEGIN
LSNRACTION=CHANGESTATUS
;
LSNRAPPLIESTO=SR
;
CLASS=SR
;
TICKETID=SRNUM
;
STATUS=INPROG
;
#MAXIMO_EMAIL_END
```

XML change status: Changing status of an existing service request

Body:

```
<MAXIMOEMAILCONTENT>
<LSNRACTION>CHANGESTATUS</LSNRACTION>
<LSNRAPPLIESTO>SR</LSNRAPPLIESTO>
<STATUS>QUEUED</STATUS>
<TICKETID>SRNUM</TICKETID>
<CLASS>SR</CLASS>
<SITEID>BEDFORD</SITEID>
</MAXIMOEMAILCONTENT>
```

7 Cron Task Setup

You can use the Cron Task Setup application to add cron tasks, cron task instances, remove cron tasks or their instances, and to modify cron task parameters. You can also change the **Active?** status or adjust the schedule of a cron task. cron tasks can be rescheduled and parameter values can be changed without stopping and restarting the server. The server performs cron tasks a specific number of times, following a schedule, and without user interaction.

Creating cron tasks requires programming resources to create custom class files.

Cron Tasks Included with the System

This set of scheduled jobs runs as part of the system server.

Name	Description
ReorderCronTask	Reorder cron task Determines the rules or parameters for scheduled reordering, direct issue, and inventory items
PMWoGenCronTask	Preventive maintenance work order generation Runs and generates scheduled work orders for planned maintenance
KPICronTask	Generates key performance indicators
LDAPSYNC	LDAP sync Synchronizes information stored in external directory servers for user authentication
ESCALATION	Escalations Escalation processes ensure that people complete critical tasks on time
LSNRCRON	E-mail Listeners Runs continuously on the system application server and processes inbound e-mail through a staging table
JMSQSEQCONSUMER	Used by the IBM® Maximo® Enterprise Adapter for polling the queue
IFACETABLECONSUMER	Used by the Enterprise Adapter for polling interface tables

Name	Description
SwSuiteCronTask	SwSuite Inspects the software titles collected in Deployed Asset, and determines whether the set of titles defined in the Deployed Asset Software Suite application are present If so, the Suite displays when inspecting that node for software discovered
ReconciliationCronTask	Reconciliation Runs reconciliation Tasks (consisting of Link and Comparison rules) to determine how assets are performing relative to the discovered data in Deployed Asset Outputs from this task: ▼ RECONLINK table that links assets to their counterpart assets ▼ Reconciliation Results table that lists the differences between compared and Deployed Assets
MeasurePointWoGenCronTask	Generates work orders when meter readings or measurements reach a condition defined in the Condition Monitoring application.
FLATFILECONS	Flat file inbound processing through cron task
XMFILECON	XML file inbound processing through cron task
VMMSYNC	Invokes IBM WebSphere® Virtual Member Manager through cron task Invokes WebSphere Virtual Member Manager APIs to populate database tables with user group and group membership records
BBCron	Periodically updates the count for the number of bulletin board postings

TIP All cron tasks are set to FULL access level, except ESCALATIONS and LSNRCRON (READONLY).

Deleting Users in the Directory Server

If you delete a user in the directory server (sometimes called the LDAP server), the VMMSYNC cron task does not delete them from the system tables. This behavior is for auditing purposes in regulated industries. If you want to delete users, you can delete them in the Users application, set up an archiving process, or create a cron task to remove them. You can use the VMMSYNC cron task to fully synchronize users and groups in the directory server and the system, and can set it up to run at a regularly scheduled interval.

Deleting Security Groups in the Directory Server

If you delete a group in the directory server, the VMMSYNC cron task does not delete it from the system tables. If you want to delete groups, you can delete them in the Security Groups application, set up an archiving process, or create a cron task to remove the groups. For information about the conditions under which you can delete a security group, see *Security Groups online help*. If you delete a group in the directory server, but do not want to delete the group from the system, you can remove all the users from that group so that they do not have access to the applications to which the group had access. To do so, complete the steps in option 1 or option 2:

- Option 1:

  1. In the directory server, remove all the users from the group that you want to delete.
  2. Wait for the VMMSYNC cron task to fully synchronize users and groups in the directory server and the system.
  3. Remove the group in the directory server.

- Option 2:

  1. Remove the group from the directory server.
  2. Go to the Security Groups application.

Remove all the users in the group that you deleted in the directory server.

Viewing Hidden Cron Tasks

READONLY tasks are hidden. You can view their parameters:

1. Go to Configuration. Select the Cron Task Setup application.
2. From the List tab, delete FULL from the **Access** field.
3. Press **Enter**. Tasks display on the List tab.

Cron Task Definitions and Instances

Cron tasks have a definition (name, class name, access level, and description).

This sample CRONTASKDEF table is populated with MAXDEMO data.

CRONTASKNAME	CLASSNAME	DESCRIPTION
ReorderCronTask	psdi.app.inventory.ReorderCron	Reorder Crontask
PMWoGenCronTask	psdi.app.pm.PMWoGenCronTask	PmWogen Crontask
KPICronTask	psdi.app,kpi.KPICron	KPI Crontask Cron Task runs KPI that are not real time
LDAPSYNC	psdi.security.ldap.LdapSyncCronTask	Synchronizes Uses and Groups from Directory Server
ESCALATION	psdi.app.escalaction.engine.EscalationCronTask	Performs Escalations
LSNRCRON	psdi.common.emailstner.EmailListenerCron	Cron task for Email listner
JMSQEQCONSUMER	psdi.iface.jms.JMSQueueCronTask	JMS Sequential Queue Consumer
IFACETABLECONSUMER	psdi.iface.intertables.ifaceTbCronTask	Interface Table Polling Task
SwSuiteCronTask	psdi.app.dpidasset.SwSuiteCronTask	Software Suite Identification Crontask
ReconciliationCronTask	psdi.app.recontask.engine.ReconCronTask	Reconciliation Crontask
MeasurePointWoGenCronTask	psdi.app.measurement.MeasurePointWoGenCronTask	Measure Point Wogen Crontask

You can create multiple instances for each definition. Each instance has an entry in the CRONTASKINSTANCE table. The attributes of the instance include:

- Set schedule string (defines the schedule for this instance)
- Description
- Flag indicating whether the instance is active
- Datetime field indicating the date and time the load/reload of the cron task is requested (not displayed to users)
- Run as User ID

  This sample CRONTASKINSTANCE table is populated with MAXDEMO data.

CRONTASKNAME	INSTANCENAME	RELOADREQTIME	SCHEDULE	ACTIVE	DESCRIPTION
LDAPSYNC	LDAPSYNC01		5M,*,*,*,*,*,*,*,*	0	Active Directory Sync
KPICronTask	KPINONREALTIME		1h*,0,*,*,*,*,*,*	0	KPI Cask Instance runs KPI that are not real
JMSQSEQCONSUMER	SEQQOUT		30s,,*,*,*,*,*,*,*,*	0	Sequential Queue Out Consumer
JMSQSEQCONSUMER	SEQIN		30s,,*,*,*,*,*,*,*,*	0	Sequential Queue In Consumer
SwSuiteCronTask	SwSuiteCronTask1		1h,*,0,*,*,*,*,*,*	0	Software Suite Identification Task
ReconciliationCronTask	ReconciliationCronTask1		1h,*,0,*,*,*,*,*,*	0	Reconciliation Task
ESCALATION	ESCALEASSTDUE	11/2/2004 1:33:09 PM	24h,*,0,*,*,*,*,*,*	0	Notify the current owner and their manager 90
KPICronTask	kpitest	11/22/2004 2:18:02 PM	1d,0,0,0,*,*,*,*,*	0	KPI CronTask. Cron Task will run KPI that are

- Instances share the same set of parameters (see the next section) but each has its own set of values and schedule. For example, the Reorder definition contains the parameter storeroom. You can modify the frequency in these instances:

- `ReorderBedford` runs daily for the central storeroom.
- `ReorderLondon` runs weekly for a remote storeroom.

See the database tables CRONTASKDEF, CRONTASKINSTANCE, and CRONTASKPARAM.

Cron Task Parameters

The cron task class file lists parameters. Parameter tables store parameter values for cron task instances.

When you create an instance, the system retrieves parameter names from the cron task class file. For each parameter, the system adds a row to the parameter table for this instance.

When instances are initialized and their parameters modified, they dynamically obtain the modifications from the database.

This sample CRONTASKPARAMETER table is populated 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	SA	directisue	
ReorderCronTask	SA	emailto	
ReorderCronTask	SA	ignoreorderpoint	0
ReorderCronTask	SA	leadtime	0
ReorderCronTask	SA	logfile	
ReorderCronTask	SA	storeroom	Central,Bedford
ReorderCronTask	SA	useagreement	1

Disabling Cron Tasks

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.

You can prohibit all or a selected set of instances from running by modifying the maximo.properties file. (in the <*Maximo root*> applications\Maximo\ properties folder).

In this example, the **ReorderCronTask01** instance of the reorder cron task is set not to run:

```
// Cron Task Manager property.
//-----------------------------------------------------------------
-----------------
//Exclude the listed cron task instances from being loaded by
this server.
//use ALL for not running any cron task.
//mxe.crontask.donotrun=ALL
//Or specify the cron task instance by crontaskname.instancename
mxe.crontask.donotrun=ReorderCronTask.ReorderCronTask01
```

If you modify maximo.properties, rebuild and redeploy the EAR file. For instructions, see *The Enterprise Application Archive Files* on page 223.

For other storerooms:

- Reorders can occur every Friday.
- Agreements are not required.
- E-mail notifications go to the supervisors in charge of each storeroom.

You can duplicate the rules of the central storeroom and modify the schedule to create the reorder cron tasks (or instances) for other storerooms.

8

Domains

About Domains

Some system fields are associated with value lists from which users select appropriate values. These lists of defined values are called domains. The system uses many domains in its applications.

You use the Domains application to add or modify domains to fit your business practices. The system uses the following domains:

- SYNONYM - These are special, reserved domains in the system. You cannot add new SYNONYM domains or delete existing ones. You can add new synonym values.

  For example, if your company procedures require two people to approve a work order, you can add synonym values for the internal WAPPR value. You can then present two different values to the user; for example, WAPPRMAN and WAPPRVP, to represent approvals at the manager and vice president level.

- ALN - A simple list of values that use one of the alphanumeric data types.

  For example, your company requires that calendar information is consistent, and you 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.

- NUMERIC - A simple list of values using one of the numeric data types.

  For example, a list containing the numbers 10, 25, 50, 75, and 100.

- NUMERIC RANGE - A list of numeric values that you define when you specify a range.

  For example, you want to track the temperature range for a piece of equipment with range values of <50, 50-59, 60-69, 70-79, 80-89, 90-99, and >100.

- TABLE - A dynamic set of values based on the values of another object.

  For 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.

- CROSSOVER - A special type of table domain in which the system brings back another value (or values) from the specified record.

  For example, you want the system to retrieve the serial number of an asset in the Assets application and insert it into a field in the Items application.

Tasks After Adding Domains

After adding domains, additional tasks might be required, depending on the domain and how you want the system to display it. You can use the Classifications application, Database Configuration application, or Application Designer application to complete these tasks.

Classifications and Domains

- You assign a domain to an attribute or use the Database Configuration application to assign the domain.
- You associate a domain with an attribute in the Attributes table window; no further configuration is needed.

Database Configuration and Domains

- You associate a domain with an attribute. Most domains also have a default value specified. If the attribute is required, a default value for the domain is also required.

  For example, 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.

Application Designer and Domains

- You modify the user interface as needed.

  For example, if you added an alphanumeric 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.

Organizations and Sites

The system uses many domains in its applications, and stores domains (default) at the System level.

If you apply domains to the Organization or Site level (by typing appropriate values in the **Organization** field and **Site** field), consider the following information:

- When you specify an Organization or Site for domain values, note where the domain is being used.

For example, in the Labor application, you use the SKILLLEVEL domain on the CRAFTSKILL object and specify both the Organization and Site values for domain values. When you access the Labor application and lookup 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.

- Leave the **Organization** field and **Site** field empty for all values (users in all Organizations and Sites can access them) or specify an Organization or Site for all values (users in the specified Organizations or Sites can access them).

If you disregard the preceding note, complicated outcomes can result. For example:

Value	Organization
GREEN	A
BLUE	B
RED	
Result	▼ Records in Organization A can only access GREEN. ▼ Records in Organization B can only access BLUE. ▼ Records in other Organizations can only access RED.

After specifying Organizations and Sites for values, records in specified Organizations and Sites no longer see values that have no Organization/Site specified.

Foreign Keys and Table Domains

Using domains, you can create a foreign key from a System level for Site-level or Organization-level objects. For example, if you want to add a new attribute for assets on the `TKTEMPLATE` object, you must:

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`.

See *Domains online help* for additional information.

9 Database Administration

To maintain database integrity, you must perform backups and other tasks on a regular basis.

Backing Up and Restoring the Database

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

- Store backups in a different location from your production database and application files.
- Schedule and regularly perform system and database backups.

You can back up any type of archive media:

Media	Description
Hard disk drive	(Recommended) Lets you restore your system quickly
Tape drive	▼ Slower, but you can keep multiple tapes of backups ▼ Typically includes backup software; see the drive software documentation
CDs, DVDs, diskettes	Limited capacity, but is useful for smaller databases, archive files, or specific executables

Types of Backups

See your database platform documentation for specific commands and procedures to perform 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 Include the following folders and any subfolders beneath them: ▼ All system product files (on the administration workstation ▼ 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 Backups	Duplicates only the databases	▼ Daily to ensure full recovery of data no more than one 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

Types of Database Backups

See your database platform documentation for specific commands and procedures to perform 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. **1** Shut down the application server and report server. **2** Perform backups. **3** 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 using the software during the backup. This process is more time-consuming, but can be useful to minimize downtime in 24-hour operations.

Restoring System and Database Backups

See your database platform documentation for specific commands and procedures to restore your database from a backup.

Before you perform a restoration procedure, test the process in a test environment, even if your backup procedure appears to be working properly.

Updating Database Statistics

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 10g documentation recommends against using ANALYZE to collect statistics for the Cost Based Optimizer; use DBMS_STATS instead.

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, which calls dbms_stats.gather_table_stats with cascade true. For example:

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

Oracle has two optimizer modes:

- Cost-based
- 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 (they do not if you never analyzed your tables). Do not set the optimizer mode to RULE.

Update Statistics (SQL Server)

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

RECOMMENDATION 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.

Updating the Database

The system includes a Maximo® database update utility, UpdateDB. Run UpdateDB under the following circumstances:

- After you install system application patches
- After you install any system options, for example, IBM® Maximo Mobile applications or IBM Maximo Industry Solutions.

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.

Updating the Maximo Database

Updating the system and database components typically involves the following steps:

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

Applying Application Patches

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

Running the UpdateDB Utility

After applying the application patch, run the UpdateDB utility.

Class files are located in the Maximo\tools\maximo\classes\psdi\script\en directory. Script files are located in the Maximo\tools\maximo\en directory.

1. From a command prompt, change the directory to:

   <root_maximo>\tools\maximo\

   For example: c:\Maximo\tools\maximo\

2. At the prompt, type **updatedb.bat** and press **Enter**.

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

Successfully completing a database patch update revises the database build version in the MAXVARS table.

Updating the Database 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. Each of these.xml files contains 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 alphabetical 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 the following message:

```
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

The system uses the a_customer.xml file to reference any system classes that have been customized. Because this 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 customizations 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 following format:

```
<?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>cust</dbscripts>
   <dbversion>V600-01</dbversion>
   <lastdbversion>V520-20</lastdbversion>
   <extensions>
      <mboset objectname='PO'>psdi.app.cust.POSet</mboset>
      <mbo objectname='PO'> psdi.app.cust.PO</mbo>
   </extensions>
</product>
```

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 `<mboset objectname>` and `<mbo objectname>` entries indicate that the purchase order classes have been extended.

Product_Description.xml

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.

The following is an example of a product description file:

```
<?xml version="1.0" encoding="UTF-8"?>
<product>
   <name>IT and Service Management Extension</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>
      <mboset objectname='objectname'>classname</mboset>
      <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>
```

E-Commerce Configuration 10

Setting Default Vendors

You can set default vendors for items that users order in Desktop Requisitions. The system stores this data at the Organization level.

For example, you can specify a default vendor for some non-stocked items. The buyer can change the vendor.

Primary vendor is another type of default vendor that you use with the reorder process in Purchase Orders and Purchase Requisitions. You specify it in the **Primary Vendor** field on the Reorder Details tab in Inventory.

If no primary vendor is specified for an item, the reorder process checks whether a default vendor is specified and takes that value if it exists, the system stores the primary vendor data at the Site/Storeroom level.

1. Open the Item Master application.
2. Display the appropriate item record.
3. Click the **Vendors** tab.
4. In the Vendors table window, click the **Details** icon for the appropriate vendor. (You can insert a new row and add another vendor.)
5. Check the **Default Vendor** field.
6. Save the record.

You can also set the default vendor in the Inventory application using the Reorder Details tab.

When a user requisitions this item in Desktop Requisitions, and the **Store Location** field is empty, then the default vendor appears in the **Vendor** field.

Autonumbering for Special Order Items

Special order items are items that you do not stock in inventory, so that they have no inventory item numbers. You must order them by description, which is generally sufficient to order and track.

You can generate item numbers for them at the Organization level.

1 Open the Organizations application.

2 Select the appropriate Organization.

3 Choose **Select Action > Purchasing Options > PO Options**.

4 In the PO Options dialog box, check the **Allow the Generation of Special Order Items?** box.

5 Click **OK**.

6 Click **Save Organization**.

Configuring Automatic Reordering

In the Inventory application, you can choose **Select Actions > Reorder** to reorder storeroom items. You can run this process automatically.

Use the Cron Task Setup application to specify the schedule and these parameters. For information about the Cron Task Setup application, see Chapter 7 *Cron Task Setup* on page 139.

Parameter	Description
ignorereorderpoint	Whether reorder cron task ignores the reorder point of the storeroom. 1 = true; 0 = false
logfile	The complete path of the log file for reorder cron task result, or stdout = system standard output, stderror = system standard error. If not specified, the file specified by mxe.msgLogFile is used.
emailto	The e-mail address where the reorder result is sent. Results for each storeroom are sent as individual e-mails. The mxe.adminEmail and mail.smtp.host properties must be specified to receive e-mail.

Parameter	Description
directissue	A list of Sites (semicolon-separated Site IDs) for which the reorder cron task processes direct issue items. If it is an empty string, direct order items are not reordered for any Site. For example: site1;site2
useagreement	Whether reorder cron task considers agreements. 1 = true; 0 = false
leadtime	The extra lead time the reorder cron task includes, in days. Default = 0.
storeroom	A list of storerooms (semicolon-separated storeroom comma Site pair) the reorder cron task processes at each run. If the property has an empty string, no storeroom is reordered. For example: site1,storeroom1;site2,storeroom2;

E-Commerce Capability

To engage in e-commerce transactions, your Organization and the supplier must be e-commerce enabled.

E-commerce suppliers have their catalog available at the IBM Corporation Operations Center or an external Web site. The supplier and buyer can approve relationships and create accounts with each other.

Users generate requisitions using the Desktop Requisitions application and the Purchase Requisitions application, which include searching and requisitioning screens.

- You can configure your Site to generate approved purchase orders from requisitions, and route them directly to e-commerce enabled suppliers.
- Route requisitions through a Workflow process to purchasing agents or other appointed individuals in your organization.

Approved purchase orders are sent to suppliers through Open Applications Group (OAG) XML transactions. Further transactions and notifications regarding the status of the purchase order are handled electronically.

Use the Security Groups application to grant Desktop Requisition or work order users access to Search Catalogs.

Buyer-initiated Transactions

Transaction	Description
Purchase Order Transaction	Sends the purchase order to the supplier and is the first transaction sent to the supplier.
Cancel Purchase Order	Sends the purchase order cancellation notice to the supplier.
Catalog Search	Searches external catalogs for items that can be added to a purchase or material requisition.

Supplier-initiated Transactions

Transaction	Description
Acknowledge Purchase Order	Confirms the supplier received the purchase order. Is the first transaction received from the supplier.
Shopping Cart	When the vendor's catalog is being searched, the shopping cart returns containing the items that were requested for the purchase or material requisition.
Advance Ship Notice (ASN)	Provides information detailing intents to transport specific quantities of items from a supplier to a single destination.
Invoice	Provides invoicing information about items shipped.

For supplier companies with which you regularly do business, create a company record in the system and complete its E-commerce Details section. See *Companies online help*.

To receive transaction e-mail notifications (such as Vendor Order Status or Advance Ship Notice), purchasing agents can create a Person record in the system. Administrative users can create this record for them.

Requisitioners can receive these notifications if they select the TRANSEMAILELECTION in the Profile Page of Desktop Requisitions.

Receiving Electronic Invoices

You can configure the ability to receive electronic invoices and to have the receipt of the invoice initiate a Workflow process.

1. Go to the Companies application and select the company record.
2. In the E-Commerce Details section of the Company tab, check the **E-Commerce Enabled** box.
3. Select the **Vendor Sends Invoice** check box.

If the **E-Commerce Enabled** box is selected, but the **Vendor Sends Invoice** box is not selected, the company must send manual invoices.

When an electronic invoice is received, it creates a record in the Invoices application, populating the INVOICE and INVOICELINE tables.

Attached Document Configuration and Administration

11

The Attached Documents action is found in most system applications. This action lets you attach relevant information to a record (or to a task on a record) in the form of a file or URL address. These attached documents can be located on your company's network, on the Internet, or in a Document Management System and can include text files, spreadsheets, and diagrams.

When you configure the Attached Documents action, 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, or 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 Administration

The Attached Documents action lets you create a document library and organize documents into folders. The Maximo® database includes the following folders that you can use for this purpose:

Folder	Contents
Attachments	Text files
Diagrams	Flow charts or part diagrams
Images	Graphic images, such as pictures of assets

You can also create more folders or organize the folders into functional categories such as permits, part sheets, photographs, procedures, drawings.

Administrators maintain the library, create folders as needed, and specify the folders available for each application. You can attach a document to a record even when the document is outside the document library.

To create a document library:

- Copy the file to the Attached Documents repository
- Specify a network path to the file, then attach the copy or the link to record

The online help provides detailed information about user procedures and administrative procedures.

Adding Document Folders

You can associate a new document folder with the application from which you created it. You must have administrator privileges to access this action.

To add document folders:

1 Open an application that has the Attached Documents action.

If an application has the Attached Documents action, it has the Attachments Library/Folders action in the Select Actions menu.

2 From the Select Action menu, select **Attachment Library/Folders > Manage Folders**.

3 Click **Add a New Document Folder**.

4 To add a new document folder, complete the fields.

Field	Description
Document Folder	Provide information about the document folder that you want to add. For example, Permits, Drawings, and Schematics.
Document Folder Description	Describe the folder.
Default File Path	Type the full path name where the files are stored. This path can be a mapped drive on a separate file server.

5 Click **OK**.

When you add a new 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.

Associating Document Folders with Applications

Associate document folders with an application before you can attach documents in those folders. You must have administrator privileges to access this action.

By default, the Attachments folder, the Images folder, and the Diagrams folder are included in every application that has the Attached Documents action.

1 Open an application that has the Attached Documents action.

If an application has the Attached Documents action, it has the Attachments Library/Folders action in the Select Actions menu.

2 From the Select Action menu, select **Attachment Library/Folders > Associate Folders**.

3 Click **New Row**.

4 In the **Document Folder** field, specify 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.

Managing the Document Library

You can store documents on 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. The following tasks describe how to manage the document library.

Adding a File Attachment or a URL to the Library

This action lets you attach a file attachment or URL to a record (or to a task on a record):

1 Open an application that has the Attached Documents action.

If an application has the Attached Documents action, it has the Attachments Library/Folders action in the Select Actions menu.

2 From the Select Action menu, select **Attachment Library/Folders > Manage Library**.

Adding a File Attachment

3 Complete one of the following steps:

a Click **Add a Document to the Library > Add New File**.

b Click **Add a Document to the Library > Add New Web Page**.

4 To add a file attachment or URL, complete the following fields:

Field	Description
Select a folder	Type a folder name or select one from the list.
Specify a file	Type the file path or browse to select the name of the file and the complete file path. Provide this information only when you are adding a file attachment.
Specify the URL	Type the URL or the global address of the page on the internet or your intranet. Provide this information only when you are adding a URL.
Name the document	Type a file name or a Web page.
Description	Type a description of the file or Web page.

Field	Description
Copy document to the default location set by your administrator (recommended)?	If you do not want the document uploaded to the network, clear this check box (which is selected by default). This field is an advanced option.
Print document with work pack?	If you do not want the document printed with a work pack, clear the check box (which is selected by default). You can enable only files with the following file extensions for printing: ▼ .pdf ▼ .xls ▼ .csv ▼ .txt ▼ .doc ▼ .gif ▼ .jpg ▼ .ppt This field is an advanced option.

5 Click **OK**.

Modifying Existing Documents

You can modify information for documents that are stored in the document library. To modify the information for stored documents:

1 Open an application that has the Attached Documents action.

If an application has the Attached Documents action, it has the Attachments Library/Folders action in the Select Actions menu.

2 From the Select Action menu, select **Attachment Library/Folders > Manage Library**.

3 In the Manage Library dialog box, click the name of the document that you want to modify.

If you click **View Details** for a document, you can only modify the **Document Description** field and the **URL/File Name** field, and select the **Print with Work Pack?** check box.

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**.

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).

To attach documents to records from within the library:

- In an application that has the Attached Documents action, click **Attachments** and select **Add from Library**.

To attach documents to records from outside the library:

1. In an application that has the Attached Documents action, click **Attachments.**
2. Select **Add New File** or select **Add New Web Page**.

   To add the document to the library, in the either the Create a File Attachment dialog box or Create a URL Attachment dialog box, you can select the **Add document to the document library for others to use?** check box.

Printing Work Packs in a UNIX Environment

To print work packs in a UNIX® environment:

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.

Attached Documents Configuration

The system uses an IBM® WebSphere Application Server or BEA® 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.

Single Computer - Windows and UNIX

For WebSphere Application Server and WebLogic Server, the single computer - Windows and UNIX scenario has the following configuration and specifications:

- The application server and the HTTP server are on the same computer.
- You 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.

Single Computer Configuration

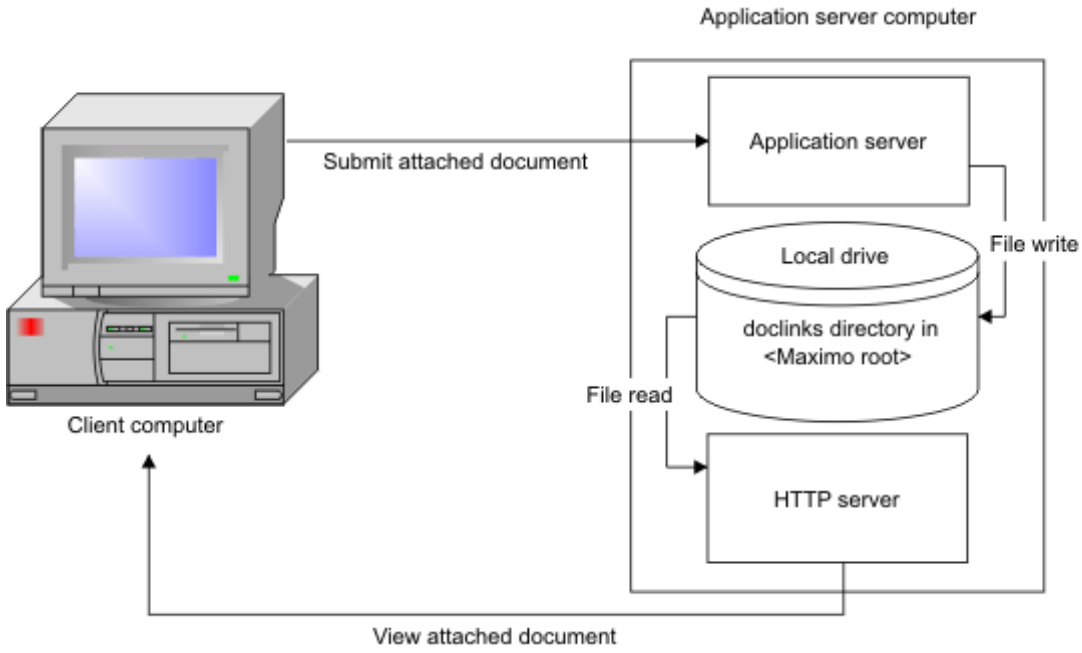

Creating Attached Documents Directories

To configure attachments, you can create an attached document directory on your computer. The following steps detail this procedure using either WebSphere Application Server or WebLogic Server.

1 Create a doclinks directory on the computer where the document files are stored. For example:

Operating System	Doclinks Directory
Windows	C:\doclinks
UNIX	/home/doclinks

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

WebSphere Steps	WebLogic Steps
5 Verify that the subdirectories were created as described in Step 3.	5 Create another directory named WEB-INF. 6 Copy the web.xml file from the deployment folder into the directory that you created (see Step 5): ▼ On Windows: <Maximo_root>\deployment ▼ On UNIX: <Maximo_root>/deployment The system contains other web.xml files. Be sure to copy the correct one. The file contains information for mapping MIME objects to customize. For more information, see *Multi-Purpose Internet Mail Extension Mappings* on page 202. 7 Verify that the subdirectories were created as described in Step 3 and Step 5.

For more information about managing document folders for attached documents, see *Adding Document Folders* on page 164.

Configuring the Application Server for Attached Documents

The following tasks described let you configure WebLogic Server or WebSphere Application Server for the attached documents action.

- If you are using a WebLogic Server, see *Create a Web Application in WebLogic* on page 170.
- If you are using WebSphere Application Server, see *Editing the httpd.conf File on the WebSphere Application Server* on page 171.

Create a Web Application in WebLogic

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 application server.
4. 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.
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 doclinks directory and select it.
9. Click **Target Module** at the bottom of the window.
10. 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.
11. 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, complete the steps described in *Create a Web Application in WebLogic* on page 170.

14 Go to *Editing Default File Paths in the System Properties Application* on page 172.

Editing the httpd.conf File on the WebSphere Application Server

In WebSphere, 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.

To edit the httpd.conf file on WebSphere Application Server:

1 Go to the location of the httpd.conf file for the IBM HTTP server. The default installation location is described in the following table:

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 Do 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, complete the steps in *Editing the httpd.conf File on the WebSphere Application Server* on page 171.

10 Restart WebSphere Application Server and the system.

Editing Default File Paths in the System Properties Application

Steps for WebSphere Application Server and WebLogic Server

After you modify the location of the doclinks directory, edit the specified file paths in the system.

To edit the specified file paths in the system, complete the following steps in the System Properties application:

1. Log in to the system. You must have authorization to edit file paths in Attached Documents.
2. Go to **System Configuration > Platform Configuration > System Properties**.
3. 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: C<PATH>\doclinks = http://<servername or IP>:<port number>/doclinks For example: C<PATH>\doclinks = http://localhost:7001/doclinks ▼ UNIX WebLogic: /home/doclinks = http://<servername or IP>:<port number>/doclinks For example: /home/doclinks = http://localhost:7001/doclinks ▼ Windows WebSphere: C<PATH>\doclinks = http://<servername or IP>/ For example: C<PATH>\doclinks = http://localhost ▼ UNIX WebSphere: /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 platform. The default value is false.	▼ Change the value to true if it is running on AIX WebSphere 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.

4 Restart the application server.

Editing Default File Paths in Related Applications

Steps for WebSphere Application Server and WebLogic Server

After you modify the location of the doclinks directory, edit the specified file paths in the system.

To edit the default file paths in related applications, complete the following steps in any application that uses Attached Documents:

1 Log in to the system. You must have authorization to edit file paths in Attached Documents.

2 Open an application that has the Attached Documents action.

If an application has the Attached Documents action, it has the Attachments Library/Folders action in the Select Actions menu.

3 From the Select Action menu, select **Attachment Library/Folders > Manage Folders**.

4 In the Manage All Document Folders dialog box, click the **Details** icon next to the document folder whose file path you want to change.

The details area is at the bottom of the page.

5 In the **Default File Path** field, edit the path to specify the new location of the associated directory. Type the full path including the mapped drive letter.

The drive letter, path, and folder names are case sensitive. They must be under the same path and folder names that you created in *Editing Default File Paths in the System Properties Application* on page 172.

Change the file paths for the Attachments, CAD, Diagrams, and Images folders to the following values:

Operating System	File Paths
Windows	C:\doclinks\attachments
	C:\doclinks\cad
	C:\doclinks\diagrams
	C:\doclinks\images
UNIX	home/doclinks/attachments
	/home/doclinks/cad
	/home/doclinks/diagrams
	/home/doclinks/images

If you create additional attached document folders, you must edit their file paths.

6 Click **OK**.

7 Restart the application server.

Changing Paths for Demo Data Library Files

Steps for WebLogic Server and WebSphere Application Server

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.

To change the file paths:

1 Log in to the system. You must have authorization to edit file paths in Attached Documents.

2 Open an application that has the Attached Documents action.

If an application has the Attached Documents action, it has the Attachments Library/Folders action in the Select Actions menu.

3 From the Select Action menu, select **Attachment Library/Folders > Manage Library**.

4 In the Manage Library dialog box, click the **Details** icon for the document whose file path you want to change.

5 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.

6 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

Operating System	File Paths
UNIX	/home/doclinks/<filename> For example, document 1001, URL/File Name is displayed as default as: \DOCLINNKS\BOILDER.DWF Change it to: /home/doclinks/BOILDER.DWF

You must change every library file path in the dialog that appears in the window.

7 Click **OK**.

8 Restart the application server.

Alternative Configurations

There are alternative configurations for the Attached Documents action when you are using either WebSphere Application Server or WebLogic Server platform.

- *Two Computers, Local HTTP Server on Windows and UNIX* on page 178:
  - 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 WebLogic Server only.

- *Two Computers, One Dedicated HTTP Server – Windows and UNIX* on page 186:
  - 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 and UNIX* on page 193:
  - 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.

Two Computers, Local HTTP Server on Windows and UNIX

WebLogic Server Only

When you use two computers, a local HTTP server, and WebLogic on either Windows or UNIX, the following configuration and specifications apply:

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).	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.
▼ You mount the Network File System that contains the document files from the document file server computer onto the application server (UNIX only).	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.

Two Computer Configuration With Local HTTP Server

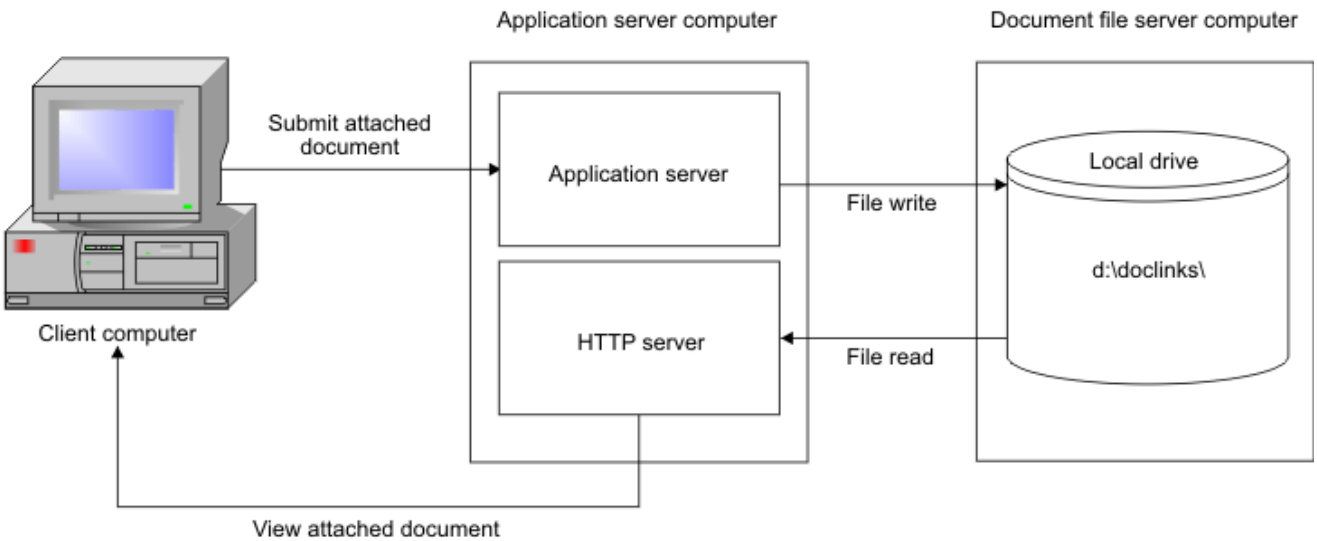

Creating Attached Documents Directories

To create directories on the computer on which the document files are stored:

For example:

Operating System	Doclinks Directory
Windows	D:\doclinks
UNIX	/home/doclinks

1. Share the drive so that users can access it.
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: <Maximo_root>\deployment
      - For UNIX, the web.xml file is: <Maximo_root>/deployment

      The file contains mime-mapping information that you can customize. For more information, see *Multi-Purpose Internet Mail Extension Mappings* on page 202.
   d. Verify the directory structure.

Creating a Web Application

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

Operating System	Path
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.
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 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, complete the steps in *Creating a Web Application* on page 179.

Editing Default File Paths in the System Properties Application

Because you modified the location of the doclinks directory, you can use the Systems Properties application to edit the specified file paths. Complete the following steps in the System Properties application:

1 Log in to the system. You must have authorization to edit file paths in the Attached Documents action.

2 Go to **System Configuration > Platform Configuration > System Properties**.

3 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. Files for such folders are uploaded to the location specified, 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 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 servername in the path must be a fully qualified servername.

4 Restart the application server.

Editing Default File Paths in Related Applications

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:

1 Log in to the system. You be authorized to edit file paths in the Attached Documents action.

2 Open an application that has the Attached Documents action.

If an application has the Attached Documents action, it has the Attachments Library/Folders action in the Select Actions menu.

3 From the Select Action menu, select **Attachment Library/Folders > Manage Folders**.

4 In the Manage All Documents Folder dialog box, click the **Details** icon next to the document folder whose file path you want to change.

5 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 in *Editing Default File Paths in the System Properties Application* on page 172.

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.

6 Click **OK**.

7 Restart the application server.

Changing Paths for Demo Data Library Files

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.

To change the file paths:

1 Log in to the system. You must have authorization to edit file paths in the Attached Documents action.

2 Open an application that has the Attached Documents action.

If an application has the Attached Documents action, it has the Attachments Library/Folders action in the Select Actions menu.

3 From the Select Action menu, select **Attachment Library/Folders > Manage Library**.

4 In the Manage Library dialog box, click the **Details** icon next to the document whose file path you want to change.

5 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 *Editing Default File Paths in the System Properties Application* on page 181.

6 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\BOILDER.DWF Change it to: H:\doclinks\BOILDER.DWF
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.DWF

You must modify every listed library file's path in the dialog box.

7 Click **OK**.

8 Restart the WebLogic Server.

Two Computers, One Dedicated HTTP Server – Windows and UNIX

The following configuration and specifications are applicable for WebSphere Application Server platform and a WebLogic server platform.

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.

Two Computer Configuration with Dedicated Document File / HTTP Server

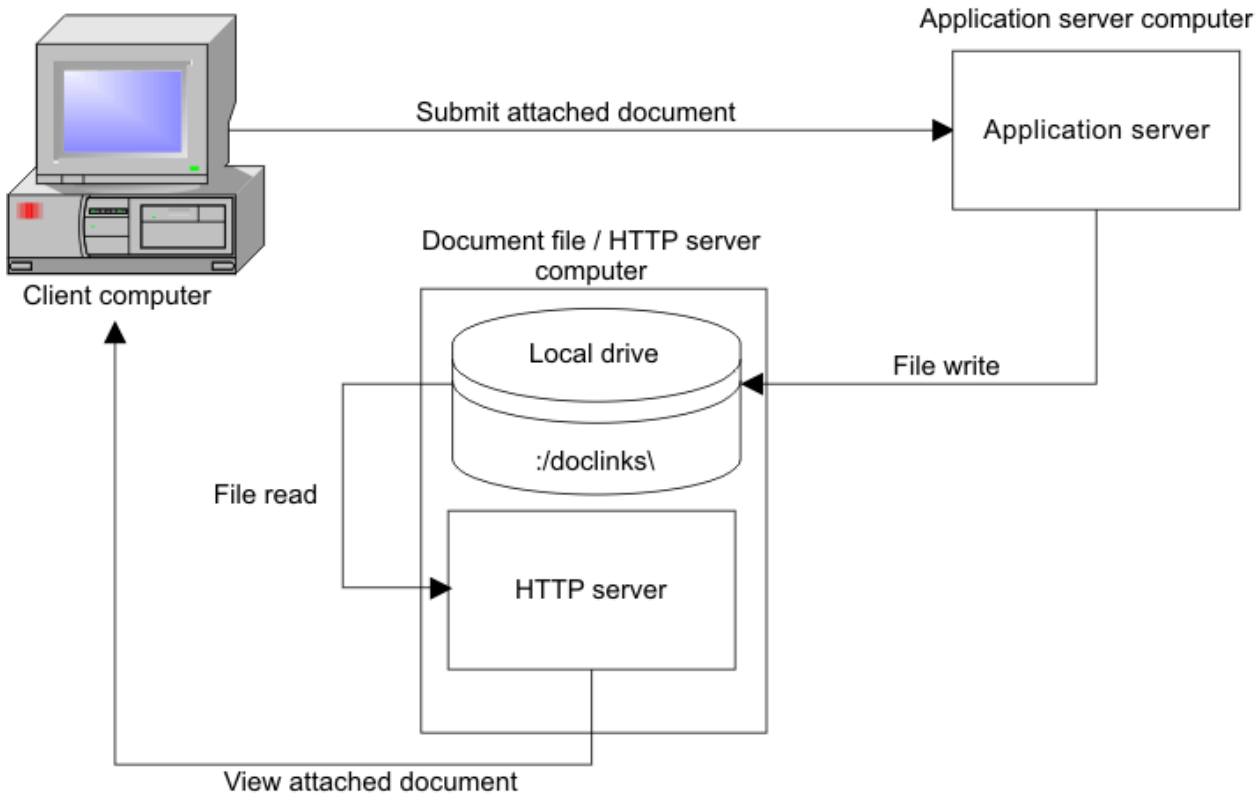

Creating Attached Documents Directories

Steps for WebSphere Application Server and WebLogic Server

To create directories on the computer on which the document files are stored:

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:

- 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 these tasks:

Operating System	Map Drive
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.

Setting up the HTTP Server for the Attached Documents Action

Steps for WebSphere Application Server and WebLogic Server

The Two Computer Dedicated HTTP Server scenario relies on an HTTP server that is independent of the system. You can use the HTTP server application that you prefer. Configuring the HTTP server for the attached documents action requires basic configuration, which is described below.

For example, to configure the HTTP server for the attached documents action:

- (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.

Because you edited the httpd.conf file, restart the HTTP server.

Editing Default File Paths in the System Properties Application

Steps for WebSphere Application Server and WebLogic 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.

To edit default file paths, complete these steps in the System Properties application:

1 Log in to the system. You must have authorization to edit file paths in the Attached Documents action.

2 Go to **System Configuration > Platform Configuration > System Properties**.

3 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. Files for such folders are uploaded to the location, mxe.doclink.doctypes.defpath.	▼ Windows: H:\doclinks ▼ UNIX: /d01/doclinks

Property	Description	Global Value
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. 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 platform. Default value is false.	▼ Change the value to true if it is running on AIX WebSphere 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.

4 Restart the application server.

Editing Default File Paths in Related Applications

Steps for WebSphere Application Server and WebLogic 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:

1. Log in to the system. You must have authorization to edit file paths in the Attached Documents action.
2. Open an application that has the Attached Documents action.

   If an application has the Attached Documents action, it has the Attachments Library/Folders action in the Select Actions menu.
3. From the Select Action menu, select **Attachment Library/Folders > Manage Folders**.
4. In the Manage All Document folders dialog box, click the **Details** icon next to the document folder whose file path you want to change.
5. 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 in *Editing Default File Paths in the System Properties Application* on page 172.

   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.
6. Click **OK**.
7. Restart the application server.

Changing Paths for Demo Data Library Files

Steps for WebSphere Application Server and WebLogic Server

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 paths to be the same as your doclinks directory setup.

To change the file paths:

1. Log in to the system. You must have authorization to edit file paths in the Attached Documents action.
2. Open an application that has the Attached Documents action.

   If an application has the Attached Documents action, it has the Attachments Library/Folders action in the Select Actions menu.
3. From the Select Action menu, select **Attachment Library/Folders > Manage Library**.
4. In the Manage Library dialog box, click the **Details** icon next to the document whose file path you want to change
5. 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 *Editing Default File Paths in the System Properties Application* on page 172.

   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.DWF
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.DWF

   You must modify every listed library file's path in the dialog.
6. Click **OK**.

7 Perform one of the following steps:

- If you are using WebSphere Application Server and you edited the httpd.conf file, restart the HTTP server, WebSphere Application Server, and the system.
- If you are using a WebLogic Server, restart the application server.

Multiple Computers, Multiple HTTP Servers – Windows and UNIX

The Multiple Computers, Multiple HTTP Servers scenario is applicable to both WebSphere Application Server platform and WebLogic Server platforms. This scenario has the following configuration and specifications:

Operating System	Configuration	Specifications
Windows	▼ You store document files on 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 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, respectively. ▼ File names are case sensitive.

Operating System	Configuration	Specifications
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.

Multiple Computers Configuration with Multiple Dedicated Document File / HTTP Servers

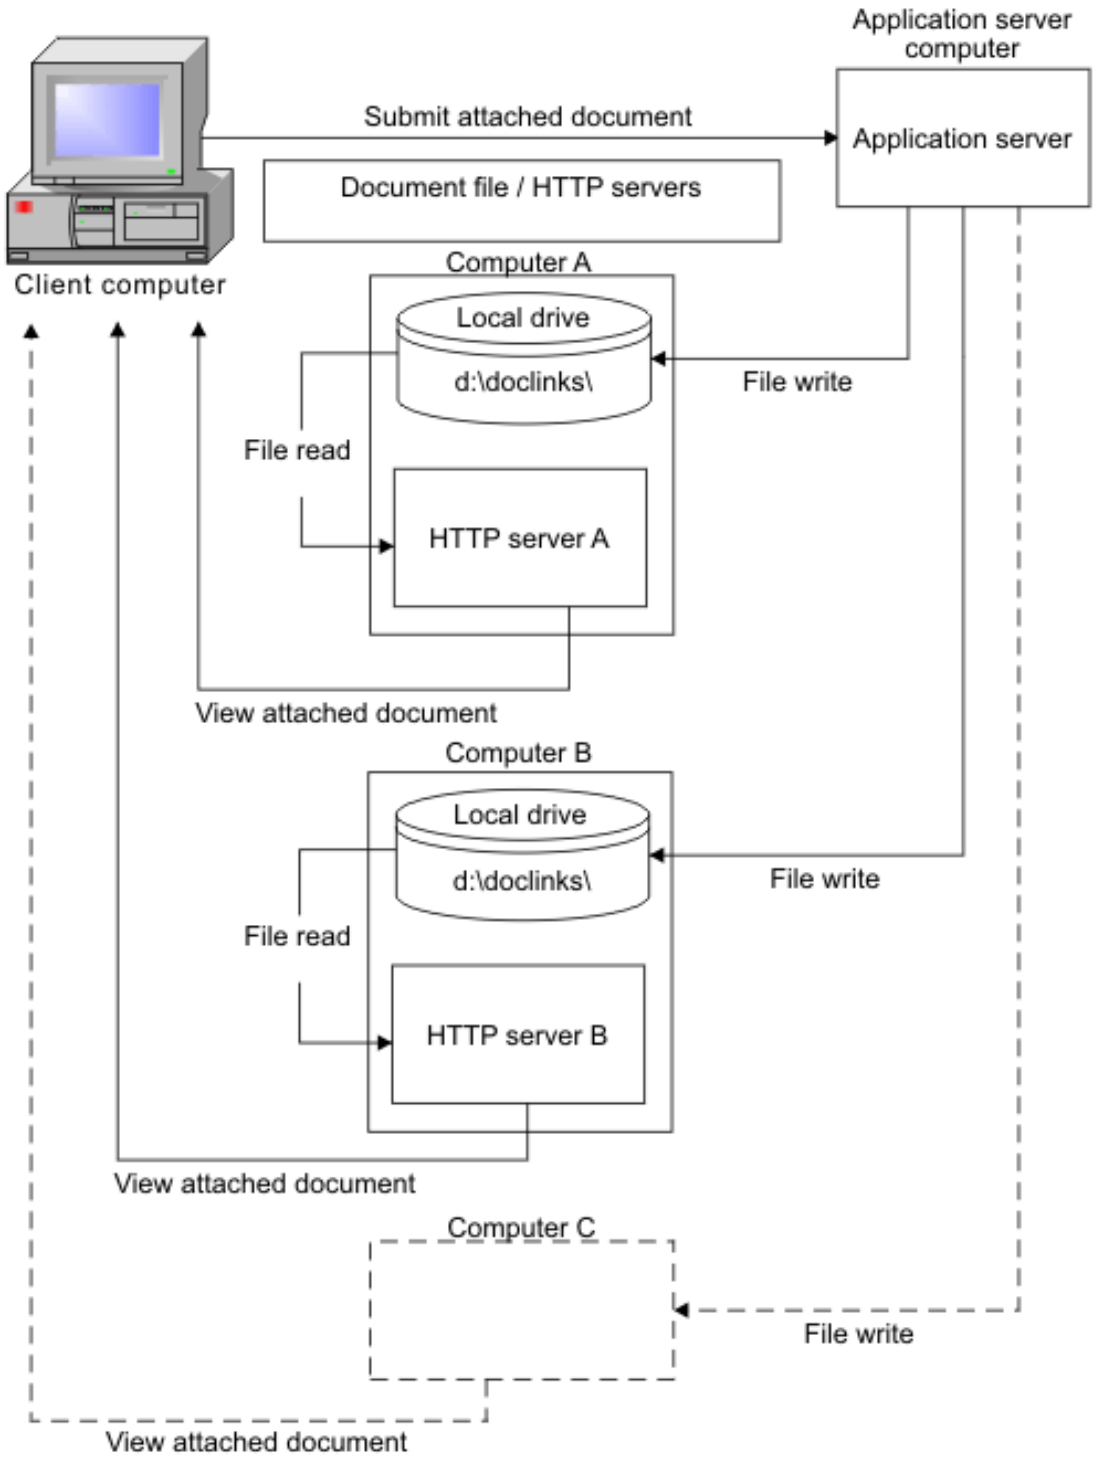

Creating Attached Documents Directories

Steps for WebSphere Application Server and WebLogic Server

To create directories on the computer on which the document files are stored:

1 Create a **doclinks** directory on the HTTP server computers that store 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 Subdirectories
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.

Setting up the HTTP Server for the Attached Documents Action

Steps for WebSphere Application Server and WebLogic Server

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. Configuring the HTTP server for the attached documents action requires basic configuration.

For example, to configure the HTTP server for the attached documents action:

- (Windows) In Apache, you edit the httpd.conf file to use d:\doclinks as its default home page documents directory.
- (Windows) In Microsoft Internet Information Services you can 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, you edit the httpd.conf file to use /home/doclinks as its default home page documents directory.

Once you edit the httpd.conf file, restart the HTTP server.

Editing Default File Paths in the System Properties Application

Steps for WebSphere Application Server and WebLogic Server

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. To edit a file path, complete these steps in the System Properties application:

1. Log in to the system. You must have authorization to edit file paths in the Attached Documents action.
2. Go to **System Configuration > Platform Configuration > System Properties**.
3. 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

Property	Description	Global Value
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 attached to 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 system copies the document 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 platform. The default value is false.	▼ Change the value to true if it is running on AIX WebSphere 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.

mxe.doclink.pathnn

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.pathnn` property, the dochost in the path must be a fully qualified server name.

4 Restart the application server.

Editing Default File Paths in Related Applications

Steps for WebSphere Application Server and WebLogic Server

Once you modify the location of the doclinks directory, you can edit the specified file paths in the system. To edit a file path, complete these steps:

1 Log in to the system. You must have authorization to edit file paths in the Attached Documents action.

2 Open an application that has the Attached Documents action.

If an application has the Attached Documents action, it has the Attachments Library/Folders action in the Select Actions menu.

3 From the Select Action menu, select **Attachment Library/Folders > Manage Folders**.

4 In the Manage All Documents Folders dialog box, click the **Details** icon next to the document folder whose file path you want to change.

5 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 in *Editing Default File Paths in the System Properties Application* on page 172.

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.

6 Click **OK**.

7 Restart the application server.

Changing Paths for Demo Data Library Files

Steps for WebSphere Application Server and WebLogic Server

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.

To change the file paths:

1. Log in to the system. You must have authorization to edit file paths in the Attached Documents action.
2. Open an application that has the Attached Documents action.

   If an application has the Attached Documents action, it has the Attachments Library/Folders action in the Select Actions menu.
3. From the Select Action menu, select **Attachment Library/Folders > Manage Library**.
4. In the Manage Library dialog box, click the **Details** icon next to the document whose file path you want to change.
5. 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 in *Editing Default File Paths in the System Properties Application* on page 172.
6. 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.DWF
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.DWF

   You must modify every listed library file's path in the dialog box.
7. Click **OK**.

8 Perform one of the following steps:

   a If you are using WebSphere Application Server and you edited the httpd.conf file, restart the HTTP server, WebSphere Application Server, and the system.

   b If you are using a WebLogic Server, restart the WebLogic Server.

Multi-Purpose Internet Mail Extension Mappings

IMPORTANT Multi-Purpose Internet Mail Extensions (MIME) mapping is only for WebLogic Server.

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.

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 Explore.

  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>`
- `<extension>`
- `doc`
- `</extension>`
- `<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 later find that you have other types of documents that do not open for viewing as attached documents, 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.

To find the MIME type for an application:

1. Access the Window registry.
2. Click **Windows Start/Run/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.

12 System Configuration

Overview of System Architecture

Access to the business components and the Web application are provided by an application server. A typical basic system configuration that you set up during installation supports a user load of 50 users or less. An advanced system configuration supports a larger user load and uses clusters of servers that can work independently and can scale up as the user load requirements increase.

The advanced system configuration information is based on the user being familiar with basic installation and setup, as well being familiar with application server features.

Basic System Configuration

A basic system configuration consists of a single instance of the system running on an application server that connects to a single instance of the database that is available on a database server. If the IBM® Maximo® Enterprise Adapter is also configured for deployment, then you must set up additional messaging queues so that the system can 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, used the advanced system configuration.

Even with fewer than 50 user loads, the basic system configuration can overload if there is a lot of 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 advanced system 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.

Basic System Configuration

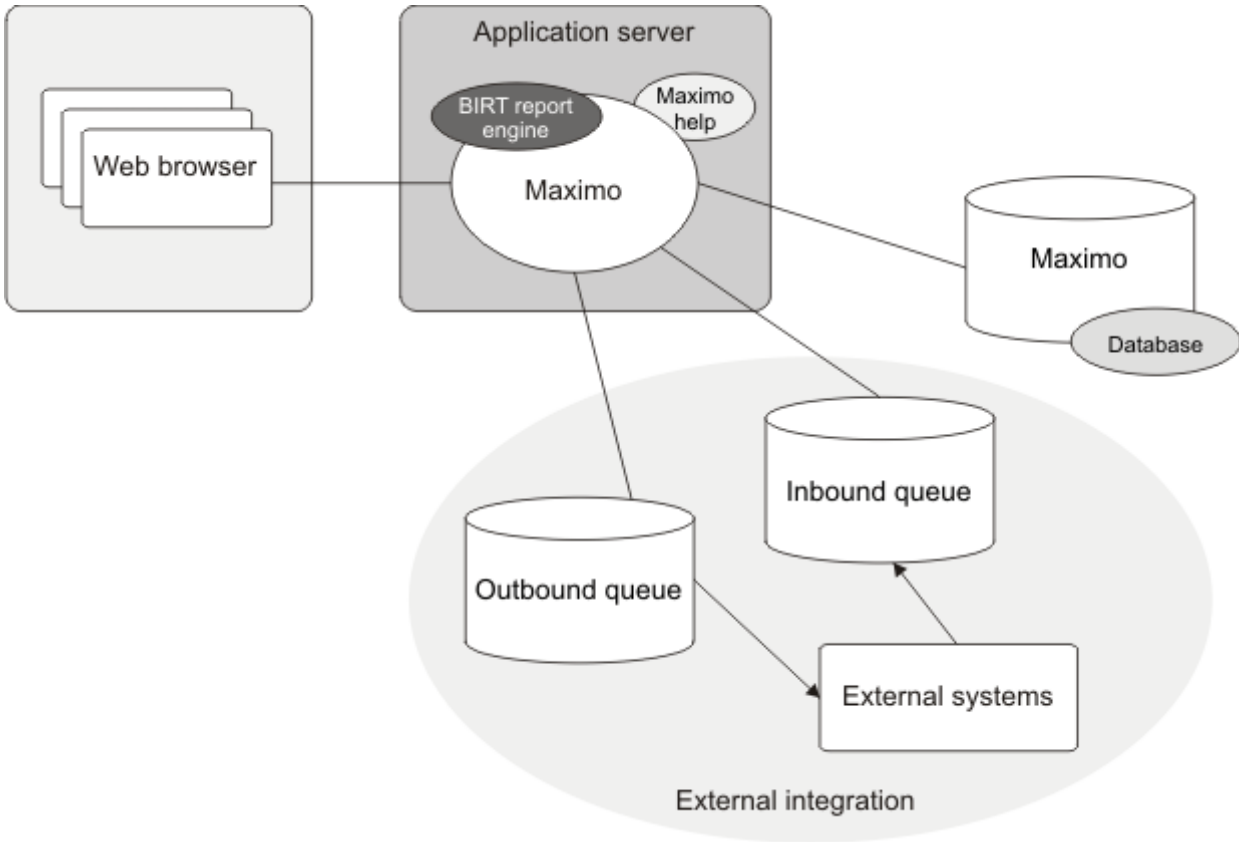

Advanced System Configuration

An advanced system configuration consists of a single instance of the system running on multiple application servers that are configured as part of various clusters. This configuration is connected to a single instance of a database that is running on a database server. Advanced configuration also provides a mechanism to set up clusters to be independent of each other.

Running on 64-bit computers, the application server process has access to more available memory than a 32-bit computer; however, the large heap space that is required can cause poor performance.

If the production system must serve a large user load, then you must implement a clustered configuration that let the system scale effectively.

When you configure the system in a clustered environment, you must configure to provide the best performance for users who are using the system from a browser. These users typically expect an immediate response from the server.

You can configure certain processes that do not require user interaction (such as scheduled jobs and inbound messages from external systems) to run in separate clusters. To enhance each cluster, you can add servers, depending on your needs.

The following diagram shows the main components of an advanced system configuration.

Advanced System Configuration

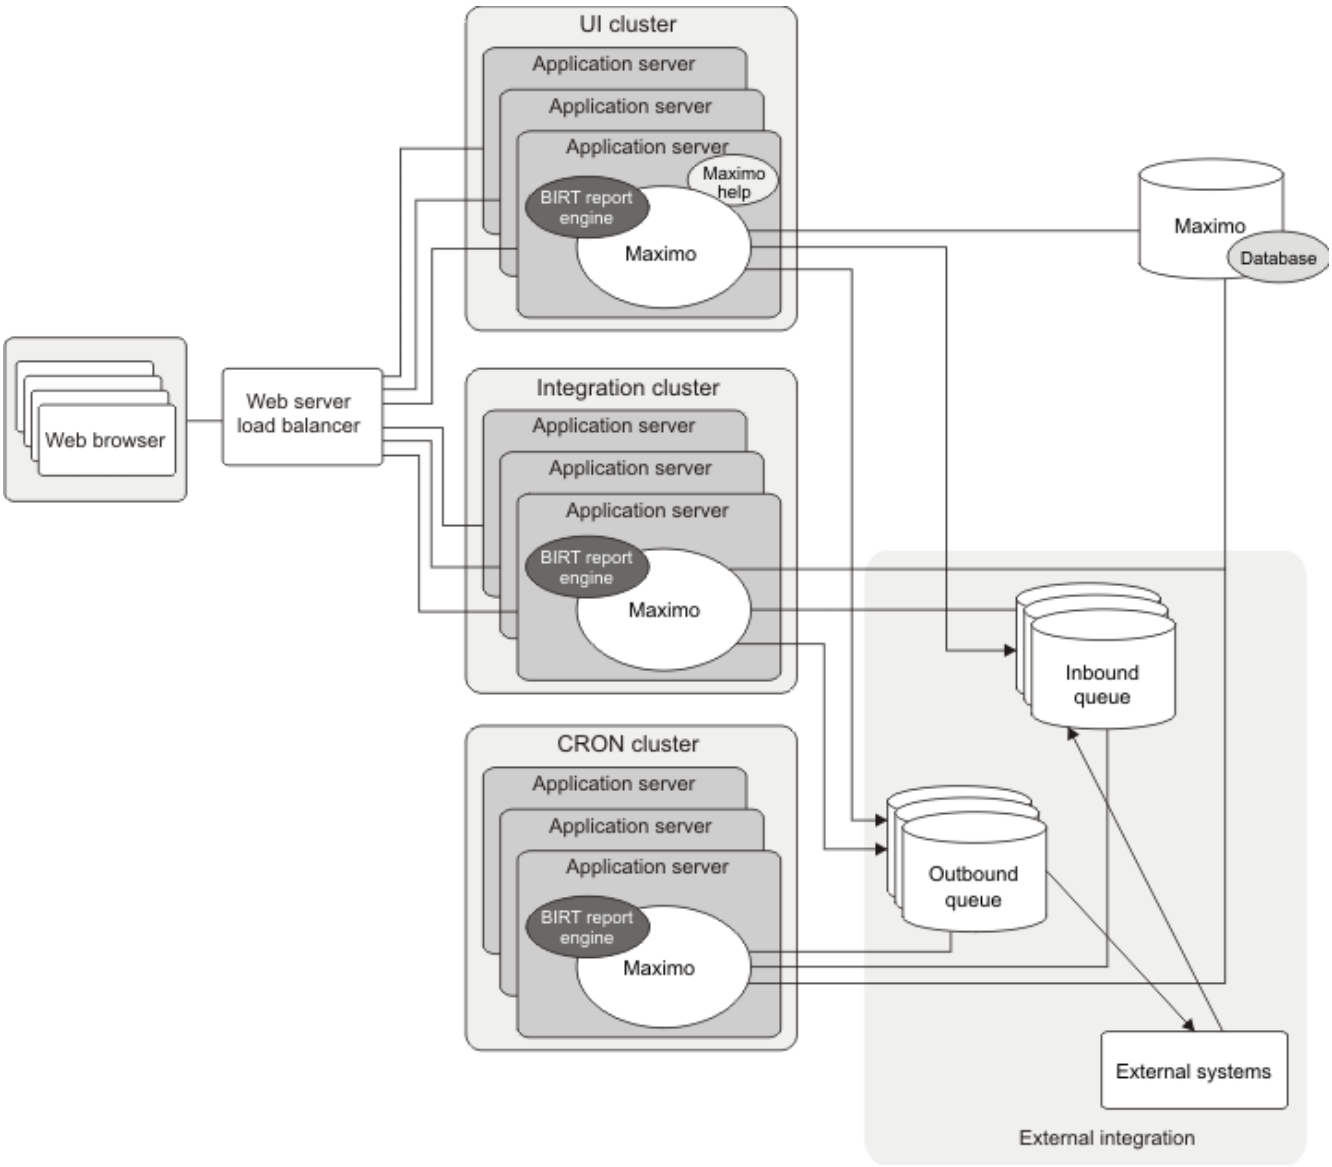

The recommended cluster configuration involves setting up three clusters:

- User interface cluster
- Integration cluster
- Cron cluster

Deploy the system in all of the cluster servers, which are used for different purposes. The default reporting engine is run from within each server that provides reporting capabilities.

- UI cluster - The UI cluster is mainly intended for users to access the system from a Web browser. This cluster is typically set up for load balancing so that the user who uses the browser only uses one URL to access the system that is running on any of the servers within this cluster.

  You can use the UI cluster and integration cluster to import and export data into the queues by using the data import/export mechanisms that are in the External System application.

RECOMMENDATION Use this feature from the UI cluster only after peak hours.

Also, transactions that users perform are sent to the *external systems* by using the outbound queues.

- Integration cluster - The integration cluster processes integration messages from message queues, and moves messages into the queues by using HTTP POST, Web services, and Enterprise JavaBeans® (EJB technology).

  You can use the integration cluster to access the system from a browser to import and export data into the queues by using the data import/export that the mechanisms that are in the External System application.

RECOMMENDATION Use the integration cluster, instead of the UI cluster, to import data. This practice ensures that users are not affected when they access the system from the UI cluster.

- Cron cluster - The cron cluster processes scheduled jobs. Some scheduled jobs can be related to integration and reports.

RECOMMENDATION The use of reports places a burden on user interactive applications because the reports are running in the same process. To avoid this problem, run reports that require system resources on separate clusters. To avoid impacting the overall performance, only run simple reports (such as work orders and purchase order printing) from user interactive applications. Configure a separate Maximo database, or an offline database for reports, to back up a copy of the production data and set up a separate system that connects to this database. With this setup, you can run reports that require significant system resources without affecting live production system performance.

Setting Up Advanced System Configuration

To set up an advanced system configuration, you must be familiar with the application server infrastructure, such as starting and stopping servers, creating and managing clusters, and maintaining advanced administration techniques for application servers. In addition, you must have the following abilities:

- Be able to set the Java™ Virtual Machine (JVM) parameter and memory settings
- Be able to set up JMS queue configurations, if integration is enabled
- Be able to build and deploy Maximo Enterprise Application Archive (EAR) files
- Be able to access the system from a browser with the appropriate URL
- Be able to set the system properties from the System Properties application

For additional information, see the link provided in *Application Server Documentation* on page 225.

The following components are included in an example of an advanced system configuration setup:

- Three clusters:
  - UI cluster - On this cluster, users interact from a Web browser.
  - Integration cluster - This cluster handles integration with external systems (for example, this cluster handles continuous queue processing).
  - Cron cluster - All scheduled cron jobs are run within this cluster, including integration-related cron jobs and scheduled reports.

To set up an advanced configuration, complete the following steps:

TIP If integration is not enabled for the system, you can omit the steps for integration.

Preparing EAR Files

To prepare EAR files, follow these steps:

1. Use the buildmaximoear command in the deployment folder where the system is installed to build two EAR files (maximo.ear), one for the UI cluster and cron cluster and the other one for the integration cluster.

   Ensure that any necessary changes to the configuration files (such as maximo.properties file) and to the deployment descriptor files (such as web.xml and ejb-jar.xml files) are made before you build the maximo.ear files.

   a. Build the maximo.ear file and rename the default deployment folder to `default-UICRON`. This maximo.ear file is deployed into the UI cluster and cron cluster.

   b. Enable integration (see the *Integration Guide* for more information about enabling integration) and build the maximo.ear file. Rename the default deployment folder to `default-INTEGRATION`. This maximo.ear file is deployed into the deployment integration cluster. These changes involve

editing deployment descriptor files to enable the message-driven beans (MDBs) for processing continuous queue messages.

IMPORTANT If you must build additional maximo.ear files because you made changes, then devise a process so that the maximo.ear files are built appropriately for the UI cluster, cron cluster, and integration cluster. (the system does not provide any scripts for such a process.) A suggested approach is to maintain only the changed files and the corresponding original files in separate folders, and write script to copy these files as needed, before you build the appropriate maximo.ear file.

2 Build the Enterprise Help Application Archive (maximohelp.ear) file, using the buildhelpear command in the deployment folder where the system is installed. This step creates the maximohelp.ear file in the default folder under the deployment folder. Rename this folder to `default-Help`.

3 Build the RMI Registry Application Archive (rmireg.war) file using the buildrmiregwar command in the deployment folder where the system is installed. This step creates the rmireg.war file in the default folder under the deployment folder. Rename this folder to `default-rmireg`.

Creating Clusters and Servers

To create clusters and servers, follow these steps:

1 Create clusters and servers:

a Create three clusters: UICluster, IntegrationCluster, and CronCluster.

b Create one or more servers under each cluster (for example, uiserver1, integrationserver1, cronserver1, and so on).

For each server, set the minimum and maximum heap sizes for the memory arguments. See the documentation for your application server for information about creating the clusters and setting heap sizes.

c Use the JVM parameter `-Dmxe.name=<name>` to name each server, where `<name>` is the name that you use to identify the server.

For example, `-Dmxe.name=uiserver1`

d Create a separate server for every computer (for example, rmiregserver1, rmiregserver2, and so on). If more than one server is set up to run on a single computer, then set up an additional server to run the system RMI Registry application. Without the system RMI Registry application, the first application server that starts creates the RMI Registry Service and all other application servers binds to the first server. If the first server fails, the remaining servers can experience problems. As a result, you need one dedicated server to host the RMI Registry service on every computer on which multiple servers are run.

RECOMMENDATION The application server that you create to run the system RMI Registry application requires a small amount of memory to start the application server. Therefore, set the maximum heap size for this application server to 128 MB.

TIP Depending on the mechanism you chose to deploy the online help (see *Online Help Configuration* on page 223), you might need to create another server, helpserver, to host online help.

Deploying in the UI Cluster

To deploy in the UI cluster, follow these steps:

1. Deploy the maximo.ear file that is in the default-UICRON folder to the UI cluster.
   a. Depending on the mechanism that you choose, use the maximohelp.ear file in the default-Help folder to deploy the maximohelp.ear file in the UI cluster or to the separate server, helpserver.
   b. During deployment on WebSphere Application Server, select the Web server, in addition to the UI cluster servers. Start the UI cluster and ensure that you can connect to the system from a browser.
2. To set up the UI cluster servers not to process any cron jobs and to disable the Report Queue Manager, go to the System Properties application and:
   - Set the `mxe.crontask.donotrun` property to **ALL** for all servers in the UI cluster.
   - Set the `mxe.report.birt.disasblequeuemanager` to **1** for all servers in the UI cluster.

   TIP If you are running multiple servers, then you must ensure that the appropriate rmiregserver that you created is started first.
3. Enable all integration cron tasks. Use the Cron Task Setup application to specify the `TARGETENABLED` property for cron task instances to a value of `1`.
4. Set the JVM parameter - `DJMSQSEQCONSUMER.SEQQOUT=1` for each server in the UI cluster on which you want the `SEQQOUT` cron tasks to run.

   This setting is necessary for the user to perform transactions that are placed in the sequential outbound queue (that was configured for the UI cluster) and sent out by this cron task.
5. Restart the UI cluster servers.

   If you are running multiple servers on one computer, deploy the rmireg application in the appropriate RIM Registry server. Ensure that the rmireg server is started first.

Deploying in the Integration Cluster

Before you deploy in the integration cluster, ensure that the JMS queue configurations are set up correctly.

To deploy in the integration cluster:

1. Access the System Properties application that is running in the UI cluster from a browser. Set up the integration cluster server not to process any cron jobs and disable the report queue manager.
   a. Set the `mxe.crontask.donotrun` property to **ALL** for all servers in the integration cluster.
   b. Set the `mxe.report.birt.disasblequeuemanager` to **1** for all servers in the integration cluster.

2. Set up the JMS queues (see *Java Messaging Service Configuration* on page 213 for the JMS queue procedure).

3. Provide the queue connection factory JNDI names:

   a. Access the External System application that is running in the UI cluster from a browser.

   b. Using the Add/Modify Queues action, provide the connection factory JNDI names that you created when you set up the JMS queues.

   TIP For BEA® WebLogic® Server, the provider URL must contain the URL of one of the integration cluster members. For example, `t3://<hostname>:<port>`, where `host name` and `port` are from one of the cluster servers.

4. Create a cron task instance, `SEQINTQOUT`, for the `JMSSEQCONSUMER` cron task:

   a. Access the Cron Task Setup application that is running in the UI Cluster from a browser.

   b. Provide the exact cron task parameters that are already defined for the `SEQQOUT` cron task instance.

   c. Set the `TARGETENABLED` parameter to 1.

5. Set the JVM parameter `-DJMSQSEQCONSUMER.SEQINTQOUT=1` for each server in the integration cluster on which you want the SEQINTQOUT cron tasks to run.

   This setting enables transactions that are caused by inbound transactions to be sent out by this cron task. The inbound transactions are placed in the sequential outbound queue, which is configured for the integration cluster.

6. Deploy the maximo.ear file in the default-INTEGRATION folder to the integration cluster.

7. Start the integration cluster. All the integration-related functions should be operational.

TIP If you are running multiple servers, ensure that the appropriate rmireg server that you created is started first.

Deploying in the Cron Cluster

To deploy in the cron cluster, follow these steps:

1. Deploy the maximo.ear file in the default-UICRON folder to the cron cluster.

2. Create a cron task instance, `SEQCRONQOUT`, for the `JMSSEQCONSUMER` cron task:

   a. Access the Cron Task Setup application that is running in the UI cluster from a web browser.

   b. Provide the same cron task parameters that are already defined for the `SEQQOUT` cron task instance.

   c. Set the value of the `TARGETENABLED` parameter to 1.

3 For integration-specific cron tasks, set the JVM parameters `-DJMSQSEQCONSUMER.SEQQOUT=1` and `-DJMSQSEQCONSUMER.SEQQIN=1` for every server in the cron cluster on which you want these cron tasks to run.

4 Enable all integration cron tasks. Use the Cron Task Setup application to specify the `TARGETENABLED` property for cron task instances to a value of 1.

   By default, only a single instance, called `SEQIN` and `SEQOUT`, is set up for the `JMSQSEQCONSUMER` cron task to read and process messages sequentially. If more than one instance must be set up, understand how instances work to sequentially process messages. See the *Integration Guide* for more information about inbound and outbound processing.

5 Start the cron cluster.

   All the cron jobs and the scheduled reports run on the cron cluster servers.

   TIP If you are running multiple servers on the same computer, ensure that the appropriate rmireg server that you created is started first.

For WebLogic Servers, if there is data that must be imported or exported into the queues through the data import/export mechanisms in the External Systems application, and if this operation must be accessible from the UI cluster, then provide the provider URL to system. To provide this URL to the system, access the External Systems application that is running in the UI cluster from a browser. Select the **Add/Modify Queues** action, and provide the provider URL for the inbound queues based on the JMS configuration that you set up. For example: `t3://<hostname>:<port>`, where `host name` and `port` are based on the values in one of the cluster member server's in the integration cluster.

Advanced system configuration is complete, and your system is running properly.

For information about cluster guidelines, see *Application Server Documentation* on page 225. See *Online Help Configuration* on page 223 for information about deploying online help. See the *Integration Guide* for information about enabling integration and build the maximo.ear file.

Java Messaging Service Configuration

The system uses JMS as the messaging standard to create, send, receive, and read messages from queues that enable 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 your application server environment and make the queues accessible through Java Naming Directory Interface (JNDI). The application servers might enforce certain limitations that are beyond the control. To set up the JMS configuration, you must be familiar with the configuration details for your application server.

Integration with external systems is supported through two message order processing mechanisms that use message queues:

- Sequential message processing - In this process, the message order is guaranteed.

- Continuous message processing - In this process, the messages are processed in parallel with MDBs, and the order in which the messages are processed is not guaranteed.

  Because the message processing order is not guaranteed, some messages that depend on a certain order can fail. For example, a vendor purchase 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 processed after the vendor record is added, the purchase order message is processed successfully.

  These errors can produce messages that remain in the queue and are picked up by another MDB and, eventually, the message is processed successfully. However, if error messages reach a specific amount for either IBM® WebSphere® Application Server or WebLogic Server, then all MDBs continuously process only the error messages. For WebSphere Application Server, the amount of error messages is equal to or greater than the maximum batch size times the number of MDBs. For WebLogic Server, the amount of error messages is equal to the maximum messages per session times the number of MDBs.

  Failed messages are not removed from the queue. Instead, you must remove the messages from the queue and replace them in the queue.

JMS Configuration for WebSphere Application Server

Application servers or clusters are members of a service integration bus. The service integration bus provides the JMS messaging support. When an application server or cluster becomes a member of a bus, the application server or bus is configured with a message engine that implements bus function of the application server (one message engine for every member). The collaboration of message engines comprises the bus.

A message engine has its own data store, which it uses to store persistent messages. A message engine supports different destinations to manage the messages:

- You configure queues based on the queue destinations.
- You set up connection factories to access the bus destinations.
- You configure an activation specification to provide queue access to the MDBs.

You set up integration with the following queues:

- Sequential inbound queue - data coming into the system from external systems that must be read in the order in which it is received
- Sequential outbound queue - data goes out of the system to external systems through this queue
- Continuous inbound queue - data coming into the system from external systems that does not must be processed in the order that t is received, and can be read 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 error handling

To address the errors that occur in the continuous message processing mode, a loop-back messaging technique is used. In this technique, the messages that were not processed in the continuous queue are placed in a separate queue, the exception queue. The messages are processed by another set of MDBs. If the messages fail again, they can be placed the same exception queue and eventually processed. For this process, set up another queue for the continuous inbound queue errors. You can design other loop-back messaging techniques or other techniques that are based on the features (and limitations) of the WebSphere Application Server that you use.

The following procedures describe how to configure JMS queues for integration.

Creating Buses

To create buses, follow these steps:

1. Create a bus named `intjmssqbus` for sequential queues. Add all three clusters (UI cluster, cron cluster, and integration cluster) as members of the `intjmssqbus`. By default, this process creates a message engines for each cluster.

2. Create a bus named `intjmscqbus` for continuous queues. Add the integration cluster and UI cluster as members of the `intjmscqbus`. This process creates message engines for the integration cluster and the UI cluster.

   The UI cluster is added to the `intjmscqbus` to provide access to the bus and to support the bulk data import operations into the continuous queue from user interface applications. The message engine that is created for the UI cluster is not used because the continuous queue messages are processed by the integration cluster.

   If necessary, you can create additional message engines for the servers in the integration cluster. For more information, see *Configuring Multiple Message Engines* on page 218.

Creating Connection Factories to Access Bus Destinations

When you create the connection factories, increase the maximum connection from 10 to 50, depending on the load.

To create connection factories, follow these steps:

1. Create a connection factory for the `intjmssqbus`, and use the following values:

   - **Name** field: `intsqconfact`
   - **JNDI Name** field: `jms/maximo/int/cf/intsqcf`
   - **Bus Name** field: `intjmssqbus`

   a. Create a connection factory at the cell scope for the `intjmscqbus`, and specify the following values:

- **Name** field: `intcqconfact`
- **JNDI Name** field: `jms/maximo/int/cf/intcqcf\`
- **Bus Name** field: `intjmscqbus`

Creating Queue Destinations

Create the queue destinations on all clusters, depending on the type of queue that you configure. For sequential outbound queues, create queues for each cluster. This setup ensures that a specific cluster downtime does not affect other clusters.

To create queue destinations, follow these steps:

RECOMMENDATION Configure other queue types (sequential inbound queue and continuous inbound queue), for the integration cluster.

1. For the UI cluster bus member, create a sequential outbound queue destination, `sqoutuibd`, for the sequential queue bus destination, `intjmssqbus`.

   If you must configure additional sequential outbound queues, create the appropriate queue destinations.

2. Create all queue destinations for the integration cluster.

   a. For the integration cluster member, create two sequential queue destinations for the sequential queue bus, `intjmssqbus`. Name the queue destinations: `sqoutintbd` and `sqinbd`.

   If you must configure additional sequential queues, create the appropriate queue destinations.

RECOMMENDATION Create additional queues for the sequential inbound purposes in the sequential queue bus, `intjmssqbus`, and for the integration cluster bus member.

   b. Create two continuous queue destinations for the continuous queue bus, `intjmscqbus`. Name the queue destinations: `cqinbd` and `cqerrb`. To manage the exceptions for the continuous queue messages for `cqinbd` and `cqerrbd`, set the exception destination to `cqerrbd`.

3. For the cron cluster bus member, create a sequential queue destination, `sqoutcrondb`, for the sequential queue bus, `intjmssqbus`.

   If you must configure additional sequential outbound queues, create the appropriate queue destinations.

Creating Queues

To create queues, follow these steps:

1. Create one queue for each destination using the default JMS Provider. If you must configure additional sequential queues, create appropriate queues. Create these additional queues for the sequential queue bus, `intjmssqbus`.

   a. For the sequential queue inbound, create the queue with the following values:

   - **Name** field: `sqin`
   - **Bus Name** field: `intjmssqbus`
   - **JNDI Name** field: `jms/maximo/int/queues/sqin`
   - **Queue Name** field: `sqinbd`

b For the sequential queue outbound that you created for the UI cluster bus member, create the queue at the UI cluster scope with the following values:

- **Name** field: `sqoutui`
- **Bus Name** field: `intjmssqbus`
- **JNDI Name** field: `jms/maximo/int/queues/sqout`
- **Queue Name** field: `sqoutuibd`

c For the continuous queue inbound that you created for the integration cluster bus member, create the queue at the integration cluster scope with the following values:

- **Name** field: `sqoutint`
- **Bus Name** field: `intjmssqbus`
- **JNDI Name** field: `jms/maximo/int/queues/sqout`
- **Queue Name** field: `sqoutintbd`

d For the sequential queue outbound that you created for the cron cluster bus member, create the queue at the cron cluster scope with the following values:

- **Name** field: `sqoutcron`
- **Bus Name** field: `intjmssqbus`
- **JNDI Name** field: `jms/maximo/int/queues/sqout`
- **Queue Name** field: `sqoutintbd`

e For the continuous queue inbound, create the queue with the following values:

- **Name** field: `cqin`
- **Bus Name** field: `intjmscqbus`
- **JNDI Name** field: `jms/maximo/int/queues/cqin`
- **Queue Name** field: `cqinbd`

f For the continuous queue inbound errors, create the queue with the following values:

- **Name** field: `cqinerr`
- **Bus Name** field: `intjmscqbus`
- **JNDI Name** field: `jms/maximo/int/queues/cqinerr`
- **Queue Name** field: `cqinerrbd`

Creating Activation Specifications

To create activation specifications, follow these steps:

1 For each continuous queue that you created, set up the activation specification at the cell scope:

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

- **Name** field: `intjmsact`
- **Bus Name** field: `intjmsact`
- **Destination Type** field: `queue`
- **Destination JNDI Name** field: `jms/maximo/int/queues/cqin`

b For the continuous queue outbound, create the activation specification with the following values:

- **Name** field: `intjmsacterr`
- **Bus Name** field: `intjmsacterr`
- **Destination Type** field: `queue`
- **Destination JNDI Name** field: `jms/maximo/int/queues/cqinerr`

Configuring Multiple Message Engines

You can add additional message engines to the integration cluster so that there is one message engine for every member server in the cluster. This setup enables all of the member servers to process messages from the continuous queues (`cqin` and `cqinerr`).

To configure message engines, follow these steps:

1 Add one or more member servers to the integration cluster. For example, `integrationserver2`, and so on.

2 Add a new message engine to the bus member integration cluster in the bus. `intjmscqbus`.

3 Set the policies to target the message engines to their preferred servers, so that each member server in the integration cluster has at least one message engine.

4 Create a policy:

a In the administrative console, select **Core Groups > Default Core Groups > Policies** to target every message engine to the appropriate server in the Integration cluster.

b Specify the following values:

- **Type** field: Select **One of N**
- **Name** field: `policy name`
- **Fail Back** field: `true`
- **Preferred Servers** field: Select the servers in the order in which you want the messaging engines to run.

c In the Match Criteria section, specify the following values:

- **Name** field: value
- **Type** field: `WSAF_SIB`
- **WSAF_SIB_MESSAGE_ENGINE** field: `Message Engine Name`

Example

Your configuration has the servers, `integrationserver1` and `integrationserver2`, in the integration cluster and two message engines. For the first server, specify the following values:

- **Type** field: Select **One of N**
- **Name** field: `intclusterme0policy`
- **Fail Back** field: `true`
- **Preferred Servers** field: Select integration server1 and integration server2

In the Match Criteria section, specify the following values:

- **Name** field: `value`
- **Type** field: `WSAF_SIB`
- **WSAF_SIB_MESSAGE_ENGINE** field: `IntegrationCluster.000-intjmscqbus`

For the second server, specify the following values:

- **Type** field: Select **One of N**
- **Name** field: `intclusterme1policy`
- **Fail Back** field: `true`
- **Preferred Servers** field: Select integration server2 and integration server1

In the Match Criteria section, enter the following values:

- **Name** field: value
- **Type** field: `WSAF_SIB`
- **WSAF_SIB_MESSAGE_ENGINE** field: `IntegrationCluster.001-intjmscqbus`

JMS Configuration for WebLogic Server

When configuring the system, you use JMS servers to:

- Manage JMS queue and topic resources (these resources are defined within JMS modules that are targeted to a specific JMS server)
- Maintain information about the store that you use for persistent messages that are received on queue destinations

JMS servers and modules provide messaging support.

If you are using WebLogic server, you cannot use the loop-back messaging technique for errors that occur in the continuous message processing mode. WebLogic server does not let you 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. As a result, only the errors in the front of the queue are processed and the remaining messages are not processed, unless the error messages are deleted or consumed. To avoid this issue, set the maximum messages per session value to `-1` for the continuous queue connection factory. This 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 memory for the process.

In this configuration, set up the following queues:

- Sequential inbound queue - data coming 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 through this queue
- Continuous inbound queue - data coming into the system from external systems that does not have to be processed in the order that it is received, and can be run in parallel by multiple MDBs

Creating Queue Stores

Create JMS stores to store the persistent messages. For sequential outbound queues, create the queue stores for each cluster. This configuration ensures that specific cluster downtime does not affect other clusters. For other queue types (sequential inbound queue and continuous inbound queue), ensure that they are configured for the integration cluster.

To create queue stores, follow these steps:

1. Create a sequential outbound queue store, and name the store: `sqoutuistore`. Target the store, `sqoutuistore`, to one of the UI cluster member servers.
2. Create the following stores for each queue that you are creating. Target the stores, `sqinstore`, `sqoutintstore`, and `cqinstore`, to one of the integration cluster member 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 store, `sqoutcronstore`, to one of the cron cluster member servers.

Creating JMS Servers

To create JMS servers, follow these steps:

1. For the sequential outbound queue, `sqoutuistore`, create a JMS server and name it: `sqoutuiserver`. Target the server to the UI cluster member server.
2. For the integration cluster, create the following JMS servers. Target the servers to the integration cluster member server.
   a. For the sequential inbound queue, create a JMS server for the `sqinstore`, and name the queue: `sqinserver`.
   b. For the sequential outbound queue, create a JMS server for the `sqoutinstore`, and name the queue: `sqoutintserver`.
   c. For the continuous inbound queue, create a JMS server for the `cqinstore`, and name the server: `cqinserver`.

   RECOMMENDATION

   For this server, set the maximum bytes to a value based on your JVM maximum heap size. It is recommended that this value be approximately 10% to 20% of the maximum heap size. This setting prevents memory errors when the producer creates messages faster than the consumer can use.
3. For the sequential outbound queue, create a JMS server for the`sqoutcronstore` and name the server: `sqoutcronserver`. Target the server, `sqoutcronserver`, to one of the cron cluster member servers.

Creating JMS Modules

To create JMS modules, follow these steps:

1. Create a JMS module for the sequential outbound queue, and name the module: `intjmssqoutuimodule`. Target `intjmssqoutuimodule` to the UI cluster.

   a. Create a sequential outbound queue for the `intjmssqoutuimodule` module with the following values:

   - **Name** field: `sqout`
   - **JNDI Name** field: `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 JMS server, `sqoutuiserver`.

   b. Create a connection factory for the queues that you created using the following values:

   - **Name** field: `intjmssqconfact`
   - **JNDI Name** field: `jms/maximo/int/cf/intsqcf`

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

   a. Create a sequential inbound queue for the `intjmssqintmodule` module with the following values:

   - **Name** field: `sqin`
   - **JNDI Name** field: `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 JMS server, `sqinserver`.

   b. Create a sequential outbound queue for the `intjmssqintmodule` module with the following values:

   - **Name** field: `sqout`
   - **JNDI Name** field: `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 JMS server, `sqoutserver`.

   c. Create a connection factory for the queues that you created using the following values:

   - **Name** field: `intjmssqconfact`
   - **JNDI Name** field: `jms/maximo/int/cf/intsqcf`

Do not create a subdeployment because the connection factory inherits the JMS module target, which is the integration 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 module, `intjmssqoutcronmodule`, to the cron cluster.

a Create a sequential outbound queue for the module, `intjmssqintmodule`, with the following values:

- **Name** field: `sqout`
- **JNDI Name** field: `jms/maximo/int/queues/sqout`

Using the default name, `sqout`, create a subdeployment for these queues you, and target it to the JMS server, `sqoutcronserver`.

b Create a connection factory for the queues that you created using the following values:

- **Name** field: `intjmssqconfact`
- **JNDI Name** field: `jms/maximo/int/cf/intsqcf`

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

Set the connection factory XA transaction to enabled.

4 Create a JMS module for the continuous queue and name it: `intjmscqmodule`. Target `intjmscqmodule` to the integration cluster.

a Create a continuous inbound queue for `intjmscqmodule` using the following values:

- **Name** field: `cqin`
- **JNDI Name** field: `jms/maximo/int/queues/cqin`

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

b Create a connection factory for the queue that you created using the following values:

- **Name** field: `intjmscqconfact`
- **JNDI Name** field: `jms/maximo/int/cf/intcqcf`

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

Set the connection factory XA transaction to enabled.

Set the Messages Maximum to `-1`.

Online Help Configuration

There are different deployment options for online help:

- You can deploy the online help application archive file (maximohelp.ear) on the same server or cluster of servers on which the EAR file (maximo.ear) file is deployed. This deployment requires no changes to the configuration files.
- Another option is to deploy the maximohelp.ear file on a separate server. All application archive (maximo.ear) file deployments can see the help that runs on the separate server.

  For this deployment, change the deployment descriptor file web.xml that is in the `applications\maximo\maximouiweb\webmodule\WEB-INF` folder in which the system is installed. In this file, edit the `helpurl` value and specify a fully qualified URL to access the help application that runs on a separate server.

  For example, if the help application is running on `myhelpserver` on port 9000, the URL is: `http://myhelpserver:9000/maximohelp`. For this deployment, edit the web.xml helpurl value as follows:

```
<env-entry>
        <description>URL of the root of MAXIMO Application Help</description>
        <env-entry-name>helpurl</env-entry-name>
        <env-entry-type>java.lang.String</env-entry-type>
        <env-entry-value>http://myhelpserver:9000/maximohelp</env-entry-value>
    </env-entry>
```

The Enterprise Application Archive Files

Enterprise Application Archive 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
- maximohelp.ear
  - Maximohelp.war

WAR files	Description
maximouiweb.war	Contains the user interface-related JavaServer Pages™ (.jsp files), Java classes, static HTML files, and static image files. The buildmaximoear.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 IBM Maximo Enterprise Adapter lets the system exchange data with other enterprise systems. Users create and maintain data in one system and the Maximo Enterprise Adapter transfers it, which eliminates duplicate processing.
maximohelp.war	Provides the online help pages. The buildhelpear.xml file has information about all the files in this module.

Building EAR Files

Windows Guidelines

To build an EAR file, follow these steps:

1. Open a Command Prompt.
2. Go to C:\Maximo\deployment.
3. Run the appropriate script:

Script	Output
buildmaximoear (Windows)	Creates a **maximo.ear** file
buildhelpear	Creates a **maximohelp.ear** file

4. These scripts take several minutes to run.

   The command prompt or terminal window displays a BUILD SUCCESSFUL line.

UNIX Guidelines

In a UNIX environment, you need a Windows computer to host the system and to build the EAR files.

The EAR files are:

maximo.ear	For the system
maximohelp.ear	For the online help application

Rebuilding EAR files

You rebuild and redeploy EAR files whenever you:

- Modify .xml files or custom class files (Maximo.ear).
- Modify html Help topics (Maximohelp.ear).
- Modify settings in the maximo.properties file (Maximo.ear).

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 IBM's 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 BEA's WebLogic Server documentation:

http://e-docs.bea.com/wls/docs92/

Miscellaneous Configuration Settings

Application Server Tuning

Tune your application server parameters based on what is recommended for your application server.

Recommended Memory Settings for the Application Server Process

The application server process in which the system is deployed must be configured with right amount of memory setting, or 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 the advanced system 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)

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

- Minimum heap size - 128 MB (-Xms512m)
- Maximum heap size- 1424 MB (-Xmx1424m)

WebSphere Application Server

If you are using a WebSphere Application Server, use the following memory settings:

- Minimum heap size - 1424 MB (-Xms1424m)
- Maximum heap size - 1424 MB (-Xmx1424m)

See *Application Server Documentation* on page 225 for more information about setting up these parameters to the JVM. It is recommended that the maximum heap size is between 1-1.5GB for optimal performance.

Load Balancing

Load balancing is the distribution of the task load across multiple instances of an application. User load comes from users who are logged in. Nonuser load comes from such things as scheduled jobs (cron tasks) and Maximo Enterprise Adapter incoming transactions.

It is optional to distribute user load and nonuser load to different application servers or clusters.

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

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.

Internet Explorer Settings

To verify that the client browser checks for the current version of the page, check your Internet Explorer settings:

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

Changing Web User Interface 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 in the following location:

<*Maximo_root*>\applications\maximo\maximouiweb\webmodule\WEB-INF\web.xml

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.

IMPORTANT Increasing the session-timeout element to a higher value consumes additional memory. It is recommended not to increase to a high value.

Logging 13

Understanding Logging

You use the Logging application to create and manage log files that contain informational, warning, or error messages about the system. Logging has three main components:

- Loggers - Loggers are components that prepare log statements to be written to console or log file. Loggers are named entities or keys; for example, log4j.logger.maximo.sql.

  Loggers form a hierarchy - A logger is defined as an ancestor of another logger if logger's name 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 - You can send logging requests to multiple destinations. An output destination is called an appender. 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.

- Layouts - 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;`

Logging enables you to create and manage log files that contain informational, warning, or error messages about the system.

Managing Appenders

You can create, modify, or delete appenders in the Logging application:

1 Choose **Select Action > Manage Appenders**.

2 In the Manage Appenders dialog box, you can create, modify, or delete an appender.

See *Logging application online help* for detailed information about this action.

The system provides you with three types appenders: Console, Rolling, and Daily Rolling, none of which you can delete. The following table contains information about these types of appenders:

Appender Types

Type	Description
Console Appender	Writes log statements to the application server console.
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 maximo.log, then the renamed file is maximo.log.1.
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. For example, if you have configured the Date Pattern attribute of your Daily Rolling Appender to `yyyy-MM-dd`, when the current file is named `maximo_scheduled.log`, the renamed file is `maximo_scheduled.log.2007-06-18`.

Log File Location

WebSphere Application Server

By default the log file is located in the following IBM® WebSphere® folder:

\IBM\WebSphere\AppServer\profiles\Custom01\maximo\logs

In this path, Custom01 is the profile name.

Also, you can specify a folder for your log files. To specify a folder for log files:

1 In the Logging application, select **Select Action > Set Logging Root Folder** action from the Select Action menu.

2 In the Set Logging Root Folder dialog box, specify a location in the **Root Logging Folder** field.

See *Logging application online help* for detailed information about this action.

If you set up a separate folder for your log files, ensure that the user account that you use to run the application server has both read and write permissions on the folder.

WebLogic Server

By default, the log file is located in the BEA® WebLogic® folder:

\BEA\92\user_projects\domains\base_domain\maximo\logs

In this path example, base_domain can be the particular WebLogic domain that you configured.

Log File Names

The corresponding log file name for an appender has a default value. For example, the name of the log file of the Rolling appender is maximo.log. However, when the file is created in the designated folder, the default file name is prefixed with the host name of the application server, as well as the server name of the system.

For example, 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, which is part of the Enterprise Archive (EAR).

Changing Logging Settings

To refresh changes to settings for the **Log Level** field and **Active?** check box, and for adding new loggers, you must perform the **Apply Settings** action in the Logging application.

Logging in Multiple Server Environment or Clustered Environment

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.

Creating and Editing Logging.properties File

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. Complete the following steps to create a logging.properties file for a specific server instance:

1. In the Logging application, select **Select Action > 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 Archive (EAR) 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. For information about deploying the EAR file, see , see Chapter 12, "System Configuration," on page 223. and search for *Enterprise Application Archive Files*.

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. Since you left a logging.properties file in the EAR, the next time you restart the application server, the settings in the file, not the database, are applied.

Changing Logging Settings

If you change logging settings in the Logging application, your changes take effect immediately when you save then and then run the **Apply Settings** action from the Select Action menu.

However, if you use separate logging.properties files on separate servers in a clustered environment, rebuild and redeploy the maximo.ear file. For information about building and deploying the maximo.ear file, see , see Chapter 12, "System Configuration," on page 223. and search for *Enterprise Application Archive Files*.

System Properties 14

System Properties

You use the System Properties application to manage system properties and the values of the system properties that various system components use. The capabilities of the System Properties application are:

- Add new system property
- Modify an existing property
- Delete an existing property
- Encrypt the value associated with the property
- Associate a domain with the property
- Specify the data type for specific properties
- Apply properties and their values into the system dynamically
- Restore original out-of-the-box values that are shipped with the system

A system property is defined as a key-value pair used at a system level.

For example, the key mail.smtp.host is a system property that is used to set up e-mail. The value for the properly is SMTP mail server.

The System Properties applications lists all system properties with their descriptions. There are two types of system properties:

- Global
- Instance

Global Properties

A global property exists only at a system-wide level, which means that the property is applicable to all of the system server instances (for example, the system application server) working with a common database, including a clustered installation.

The key characteristics of global properties are:

- Online Changes Allowed - This characteristic determines whether the user is allowed to change the value of a property using the System Properties application. By default, a user can change the majority of the properties using the System Properties application. Some properties, such as mxe.db.user, cannot be changed in the application. Instead, edit the corresponding properties file, rebuild, and redeploy the enterprise archive (EAR) files for the changes to take effect.

- Live Refresh - This characteristic determines whether a new value for a system property can be applied and take effect immediately throughout the system. By default, the majority of the properties allow Live Refresh. Live Refresh is an action item available from the Select Action menu or the application toolbar. Run this action for any property changes to take effect.

Instance Properties

An instance property is defined and associated with a specific system server. For example, you can configure the system property, `mxe.crontask.donotrun`, to be an instance-specific property. To perform this configuration, access the System Properties application. In the Instance Properties table window, associate the `mxe.crontask.donotrun` system property with a specific server (for example, MXServer1) and a value applicable only to that server (for example, a value of BBCron). As a result, the Bulletin Board cron task (BBCron) does not run on MXServer1; however, BBCron can run in another system instance, such as MXServer2.

When you create both a global value and an instance-specific value for the same property, the instance-specific value takes precedence.

Encryption of Property Values

Encryption of property values determines whether the value of the property must be encrypted. The value is encrypted using the system's standard encryption function.

Security Level

Security level determines whether a property can be retrieved an authenticated connection to the system is established. There are three levels of security:

- 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 within the business object framework of the system.

Values in Properties File vs. Application

When you assign a value to a property in both a file (for example, maximo.properties) and the System Properties application, the value in the file takes precedence. This feature is useful in a development environment. For example, multiple developers can use a common database, but run separate system instances with different property values.

This feature is also useful when you want one server in a cluster to handle the reorder cron task. In this case, you can create a maximo.properties file specifically for that server instance. For more information about the maximo.properties file, see *Maximo.properties File* on page 235.

In the System Properties application, if a property is defined to be available only from the maximo.properties file, but the property is not present in the file, the application server does not start and a message is written to the maximo.log. If the **File Override?** check box is selected, the property was defined in maximo.properties 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 and a warning is written to the maximo.log.

Maximo.properties File

The following table lists properties that you must define in the maximo.properties file. If you do not define these properties in the maximo.properties file, the application server does not start and an error message is written to the log file of the system or the application server console.

Property	Description
mxe.name	Name to bind the MXServer server object to in the RMI registry.
mxe.rmi.port	RMI communication port. If set at zero, RMI uses any available port. You can select another available port number.
mxe.db.schemaowner	Owner of the database schema.
mxe.db.driver	Specification of JDBC Driver.
mxe.db.url	JDBC URL of database
mxe.db.user	Database user that the server uses to attach to the database server.
mxe.db.password	Password for the database user name.

Some properties contain password information. If you change a password, update the associated property value. These properties include:

- 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 Maximo database password, update the mxe.db.password property to reflect the change.

You can also add additional properties to this file for a specific server instance, as described in *Values in Properties File vs. Application* on page 234.

The maximo.properties file contains database-specific examples, where necessary, and is located in the following folder:

```
<Product_root>\applications\maximo\properties
```

Attached Document Properties

You use the Attached Documents action to attach various documents to individual records. For more information about using attached documents, see Chapter 11, *Attached Document Configuration and Administration*.

If you use the Attached Documents action, use the Systems Properties application to configure the following attached document properties:

Property Name	Example Property Value
mxe.doclink.maxfilesize	Replace 10 with the desired value in MB (20 = 20 MB).
mxe.doclink.doctypes.defpath	C:\DOCLINKS\
mxe.doclink.pathnn (Where nn starts at 01- by default you can configure up to ten mxe.doclink.pathnn properties in your system environment. The system has ten properties by default; however, you can add additional properties, if necessary.)	C:\Doclinks=http:// documentserver_name or IP address/

In the mxe.doclink.pathnn. properties, the documentserver_name in the path must be a fully qualified server name.

15 Bulletin Board

The Bulletin Board application an electronic message board that lets users with the appropriate security authorizations to:

- Create and view messages, in an electronic board, regarding critical problems and incidents
- Create a communication to broadcast information throughout the enterprise
- Create a message based on the Bulletin Board content, including the audience

You can specify the date and time for messages to appear on the Bulletin Board, as well as a date and time for deletion. You can also define a date and time when you want the message to be removed from the Bulletin Board.

Viewing Messages

You can view messages from the following locations:

Start Center	By default, the Bulletin Board area displays unread Bulletin Board messages.
Any system application	The Bulletin Board icon (in the navigation bar of each application) indicates whether you have messages.

You can click a message to display or collapse its details. A list of messages displays, with the most recent message at the top, including the date and time they were posted.

For additional information, see *Bulletin Board online help*.

16 Sets and Organizations

Sets

The Sets application lets you create a framework for sharing item and company (vendor) data across multiple organizations. The system stores both item and company master records in sets. These sets exist above the organization level so that organizations can share the same data.

The system:

- Stores Sets at the database level
- Associates item and company records with a category called Sets
- Stores item and company records at the Organization level (Organizations are associated with Sets)

Type	Description
Item Set	Lets Organizations choose from a common set of items. Unique identifiers are required: ▼ For each Item Set ▼ For each item in the Set Items that you create are cataloged into the Item Set associated with the same Organization to which your default insert Site belongs.
Company Set	Ensures that all Sites and Organizations use consistent names for vendor businesses. To negoiate the best prices with vendors, lets you consolidate vendor reporting and share pricing information when purchasing products or services.

For additional information, see the *Multisite Administrator Guide* and *Sets online help*.

Organizations

You use the Organizations application to set up the organizations and sites to use with your implementation. Define at least one organization and one site.

For each organization, specify some properties:

- Base currency
- Item and company sets
- Default item status
- General ledger clearing account
- Address codes for organizational units, including sites, that you can use on records

By configuring the system with multiple organizations and sites, your company can run the system in multiple facilities all from the same database. Each site can access its own site-specific data and use data that is common across sites. Rather than installing multiple instances of the system, each with its own database, a company can install the system once, with different sites all accessing the same database through a Web Browser.

The system's architecture enables different organizations and sites to keep some parts of their operations separate, while sharing others. For example, different facilities can share purchase agreements and vendors, and share and transfer items, while keeping work orders and job plans separate.

Organizations and Sites

Organizations and sites are logical divisions of a company determined by what types of operations are performed at different locations, and what data can be shared among them. An organization is a major division of a company that contains one or more sites.

A site is a subdivision of an organization that can track inventory and other data separately from other sites. Certain types of information are unique to a site and not visible to other sites within the organization. Sites belonging to the same organization must use the same currency and share the same options for work orders, assets, labor, and certain other types of data.

A site is a subdivision of an organization that can track inventory and other data separately from other sites. Certain types of information are unique to a site and not visible to other sites within the organization.

Sets and Organizations

When working with Sets and Organizations, there are some general guidelines:

- Before creating an Organization, create at least one Item Set and one Company Set.
- You must associate each Organization with one Company Set and one Item Set.
- You can create an unlimited number of Sets.
- Multiple Organizations can use the same Item Set or Company Set.

Application Levels and Data Storage

The system stores application data at one of the following four levels:

- System level - the data is available to all organizations and sites.
- Set level - a special category by which multiple organizations can share items and vendor company data, data that the system stores 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.

Applications are considered system-, set-, organization-, or site-level depending on the level at which the system stores the data. In implementing a multisite installation of the system, it is useful to keep in mind which applications are at which level, because the data storage level affects how you use the system, as illustrated in the following examples:

- It 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 for identifying 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.

- It 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 that application's records for that organization. To use a specific example, if a user at Site 1 in Organization A creates a General Ledger account in Chart of Accounts an organization level application then users at Site 2 in Organization A can also access that General Ledger account.
- It affects some user default settings. For example, if a site is used as a filter for displaying records, but the application is organization level, then a user can access all records for the organization owning that site.

Application Options

Within the Organizations application you set defaults for a wide variety of options relating to applications or groups of applications, such as work orders, inventory, purchasing, preventive maintenance, assets, and so forth. You access these options from the Select Action menu (see *Organizations online help*, and you can also access the online help from the associated dialog boxes).

In the various dialog boxes you use to set application options, you sometimes have to select a site. If you have to select a site, the setting applies at the site level. If you do not select a site, the setting applies at the organization level, except for system and set level autonumbering, and for settings in the System Settings dialog box.

17 Calendars

You can use the Calendars application to create and modify calendars associated with these system records:

- Assets
- Labor
- Locations
- Organizations
- People
- Preventive maintenance records
- Service Level Agreements
- Tools
- Work orders

Calendar records incorporate start and end dates, shift definitions, and non-working time. Holidays are examples of non-working time. Any number of person, asset, or other records can reference a single calendar.

A calendar record is defined by a start date and end date, and by the shift definitions and nonworking time you apply to it. Nonworking time includes holidays and any other type of nonworking time you want to define. Applying shifts and nonworking time to a calendar generates the work periods for the calendar.

Typically you create calendars for Organizations, but you can also make them Site-specific. You might need multiple calendar definitions. For example:

- Corporate Calendar – includes standard shifts and holidays
- Asset Calendar – working time calendar for asset UPTIME

For additional information, see *Calendars online help*.

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. Use the following applications and icons to enter exceptions to the standard calendar:

Application	Icon
People	Modify Person Availability
Assignment Manager	Modify Availability

The system combines the standard calendar assignments and the exceptions to determine a person's availability for a given day, shift, and so on.

Shift Patterns

A shift defines working time without being date-specific. You select the working days for the week, then designate the start time and end time for work. For example, create a shift called First, with these properties:

- Working days are Monday through Friday
- Work starts at 7:00 a.m.
- Work ends at 3:00 p.m.
- Work hours for the day total 8

You can create special shift definitions that are atypical for your work Site (for example, a Saturday night or Holiday shift).

Once you define a shift, you can apply it to a calendar. After you create a calendar, you can use it on person, location, asset, and other records to specify working time.

Sample Shift Patterns	Start Day
Seven days	Sunday
Multiple of seven days For example, 14, 21, and so on	Monday
Five days	Rotates

If the number is not a multiple of 7, the pattern does not repeat on the same days of the week.

18 Classifications

You use the Classifications application to create detailed information about objects that can be classified and retrieved later. These objects include assets, locations, items, tickets, work orders, solutions, configuration items, purchasing documents, and job plans.

TIP Base your classification structure on how you currently group things in your business.

Before Creating Classifications

Before you create a classification, determine the information that you want to retrieve. You must group information so that you can perform statistical analysis later. For example, to determine how many customers complained about problem A versus problem B, classify problem A differently than problem B.

Only create classifications if you intend to use them to retrieve information.

RECOMMENDATION Organize information into top-level categories, such as:

- Information Technology assets
- Production assets
- Facility assets
- Fleet assets

Work from the top levels into the more detailed levels. For example, fleet assets can contain 18-wheel trucks and sales fleet cars, or you can categorize by maintenance groups.

You can use a visual tree control to classify items and search for classified items. You can create unlimited classification levels.

TIP Build classifications top down, from parent to child levels. Do not create entire branches. Instead, work from the top levels into the more detailed levels.

Classification Standards

The system does not provide standard classifications. However, you can apply industry standards when you create classifications.

For example, standard Vehicle Maintenance Reporting Standard codes exist for vehicle or fleet maintenance, and most mechanic know that an oil change is code 42-3-2.

Base your classifications for information technology assets on the United Nations Standard Product and Services Classification (UNSPC) codes.

Contact the IBM® professional services group or industry solutions group regarding classification standards.

Using Classifications

Users can search classification structures and attributes with associated values when they use any object that can be classified, such as assets, items, locations, configuration items, work orders, tickets, and so on.

Entity	Description
Classifications	Define at the System, Site, or Organization level
Items	Define at the System level only

Service Management Examples

Service management involves the maintenance, tracking, and resolution of Service Requests that are related to the hardware and software issues with computer systems.

Example1

A service desk organization creates a four-level classification structure to categorize tickets and work orders. This classification structure helps diagnose issues, categorize work orders, enhance reporting, and other activities.

Example 2

A user contacts the service desk, requesting a Windows XP installation. A service request and change record are created. The change record is classified as `NEW SW REQUEST>OPERATING SYSTEM>WINDOWS XP`.

Defining Classifications

You must define classifications so that the system search capability can find them.

Classification Structure

The structure consists of parent/child relationships between individual nodes.

Associate Classifications

You can define associations between actual configuration item classifications and authorized configuration item classifications. Defining, or mapping, these associations lets you create an authorized configuration item from a configuration item record, and lets you include the limited number of attributes that you need for configuration management and change control.

For more information about associating classifications, see *Classifications online help*.

What You Can Classify

You can use attributes to search for objects that can be classified, such as assets, locations, items, tickets, work orders, solutions, configuration items, purchasing documents, and job plans (for example, you can search for a blue car).

See the `Synonym Domain Classusewith` in the Domains application for a complete list of objects that you can classify. To find this information, from the Start Center of the system, go to **System Configuration > Platform Configuration > Domains**.

Attributes

Each classification node contains a list of attributes (characteristics of a classification object). Associate the attributes with the truck, as described in the following table.

Classification	Attributes
Truck	▼ Horsepower ▼ Tire size ▼ Exterior color

Sections

You can break attributes into sections. Sections are groupings of attributes, allowing the same attribute to be used multiple times.

Example

You define a pipe in the system as an asset. The pipe:

- Is 80 feet (~25 m) long
- Contains 10 sections of equivalent length
- Has an interior diameter that tapers from one end to the other

  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, so the attribute is the interior diameter.

Integrating Classifications with Other System Applications

You can integrate classification with other applications in the system. You can create classifications for activities, assets, changes, incidents, items, locations, problems, releases, service requests, solutions, work orders, configuration items, purchasing documents, and job plans. (See the Synonym `Domain Classusewith` in the Domains application for a complete list of objects that can be classified. To find this information, from the Start Center of the system, go to **System Configuration > Platform Configuration > Domains**).

You can search for classifications:

- Use any application that contains an **Asset** field**, Location** field**, Item** field**, Ticket** field**, CI** field**, Job Plan** field**, Work Order** field, or **Solution** field.

  For example, from a work order, you can search for pump-related items or toner cartridge-related solutions.

- On the List tab, use the advanced search option in the following applications:

  - Activities
  - Assets
  - Changes
  - Incidents
  - Item Master
  - Locations
  - Problems
  - Releases
  - Service Requests
  - Solutions
  - Work Order Tracking
  - Configuration Items

  For example, you can use the List tab to generate a results set of work order records based on a classification.

For additional information, see *Classifications online help*.

19 Chart of Accounts

The Chart of Accounts application lets you set up default general ledger accounts and resource codes for some standard accounting functions. You can use this application to:

- Create general ledger account codes and components
- Define financial periods
- Create default general ledger accounts

Chart of Accounts online help contains information about using menu items, and other topics not included in this chapter.

Tasks in Other Applications

To use the chart of accounts functions appropriately, specify information in other system applications:

- In the Database Configuration application, specify the format of general ledger account codes using the **GL Account Configuration** action. The format includes the number and length of components, delimiters, and so on.
- In the Organization application, define tax codes, rates, and dates using **Purchasing Options > Tax Options**.

See the *Finance Manager Guide* for additional information about general ledger database columns by application and table.

General Ledger Account Codes Overview

General ledger account codes typically include components (segments) separated by delimiters.

Example

```
6000-200-350
```

Placeholder characters represent components without values.

Example

```
6000-???-350
```

You define the format of the code in the system:

Application	Description
Database Configuration > GL Account Configuration	Define the number, length, and data type of components, and whether the components are required, and delimiters
Organizations > System Settings	Specify placeholder characters

You can use the Chart of Accounts application to define the value of component values, then link component values to create general ledger account codes for financial tracking.

You can also specify the validation rules for component combinations for the general ledger account codes that users can use. Users can use:

- A combination of any existing components
- Only component combinations that have been specified as general ledger accounts

Standard Accounting Functions

This application lets you create default general ledger accounts and resource codes for many standard accounting functions. You typically create accounts and resource codes within the system to correspond with accounts you use in your external accounting system.

See *Chart of Accounts online help* for information about general ledger accounts.

Merging General Ledger Accounts

In some instances, a general ledger account field might not be uniquely specified (for example: a general ledger account for a location and general ledger account for an asset). Generating transactions, such as work orders, typically requires choosing defined account component values, and the system invokes a set of rules to handle them.

General ledger accounts are merged component by component. Defined components supersede undefined components. Suppose the first components of two account codes are `6000` and `????`; the merged first component is `6000`.

See *Chart of Accounts online help* for information about general ledger accounts.

Working with General Ledger Accounts

You can add or modify general ledger accounts and account components, and create default accounts.

Typically, general ledger accounts are downloaded from the general ledger chart of accounts established in your accounting system. You can also create general ledger accounts in the system, at the Organization level. Each Organization has its own chart of accounts definition.

Deactivating Values

If you deactivate a component value, all general ledger account codes with that component value become inactive.

For example, the value of an active Resource component is 888. You deactivate component value 888, close Chart of Accounts, and reopen it. The general ledger Accounts that use 888 for the Resource component are 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. For example, a work order uses a cost center value of 6250, which you deactivate. The work order still uses cost center value 6250.

Reactivating Values

If you reactivate Resource component 888, the system asks whether to reactivate the corresponding account codes.

Downloading Account Codes from an 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®. You purchase product-specific application programming interfaces separately.

See your IBM sales representative. Creating your own application programming interface for your financial system is possible.

Cost Management 20

The Cost Management application supports the integration of the enterprise application management system and Oracle® Projects applications. If the Adapter for Oracle application is installed and the Projects interface is activated, the Cost Management application can retain projects, tasks, and cost summary information from Oracle Projects.

You can use this application for other purposes; however, no business rules or processes are defined for other uses such as cost rollup. The Cost Management application does not offer standard reports.

For more information about cost management, see *Cost Management online help*.

21 Currency Codes

You use the Currency Codes application to define currency codes and specify which codes to be used. A currency code is a short, user-defined value that you create to represent a currency; for example, CND for the Canadian dollar.

The Currency codes application stores currency codes at the system level. All organizations can view and use the defined currency codes, and add new ones as needed. First create a currency code, then specify the base currency and set up the exchange rate.

After you establish an active currency code, you can:

- Use the currency code wherever **Currency** fields appear (purchase requisitions, purchase orders, invoices, and companies)
- Use other applications for currency administration:
  - **Organizations**: specifies an Organization's base currency.
  - **Exchange Rates**: specifies exchange rates between currencies

For more information about currency codes, see *Currency Codes online help*.

22 Exchange Rates

The Exchange Rates application lets you type, view, and modify exchange rates that are used to convert currencies.

When a user types an amount in a foreign currency, the active exchange rate are located and the cost in the base currency of your company is calculated. If the Exchange Rate application finds that the exchange rate between two currencies is not defined, it uses specific rules and logic to calculate the exchange rate from other exchange rates, if they exist.

Exchange rates are stored at the Organization level. Therefore, each Organization defines its own exchange rates. Currency codes are stored at the System level and are available to all Organizations.

You can use other applications for currency administration:

- Currency Codes - defines currency codes
- Organizations - specifies an Organization's base currency

Rules and Logic

Two Currencies

Defining exchange rates implies inverse relationships. When defined rates are not found for a given date, the system checks whether inverse relationships are defined and uses them to calculate rates.

For example, if the rate from currency A to B is 4.0, then the rate from currency B to A is 0.25 (if 1 A = 4 B, then 1 B = 0.25 A).

If you specify only an A to B rate, and the cost of a purchase order item is in currency B, users can specify currency A on a purchase order and the system converts to currency B.

Three Currencies

When two currencies are independently defined relative to a third, rates can be calculated.

For example, with these defined rates:

- A to C
- B and C

The system can calculate:

- A to B
- B to A

If 1 A = 2 C and 1 B = 4 C, then B is twice the value of A.

Therefore, 1 B = 2 A and 1 A = 0.5 B:

Relationship	Value
A to C	2.0
B to C	4.0
A to B	0.50
B to A	2.0

One currency must be the Base 1 currency.

Properties

- You can define multiple rates between the same two currencies (A to B, for example). Dates cannot overlap.
- On any given date, you can define only one exchange rate between two currencies.
- If there is a gap between rate periods for a currency pair, for example, a month when no rate is specified, the system finds no exchange rate.
- The expiration date of an exchange rate expires at the first millisecond that the exchange rate is no longer valid.

Converting Foreign Currencies to Base Currencies

When you specify a foreign currency on a purchase requisition or purchase order, three values are calculated:

- Total Cost, expressed in the foreign currency
- Total Base Cost, expressed in the base currency of your company (Base 1)
- Total Base 2 Cost, expressed in the second base currency of your company (Base 2)

Example

Your base currency is USD. Your second base currency is CND. You type a purchase requisition for gaskets, to be ordered from a French company (using Euros).

1 On the PR tab, select EUR (Euros) in the **Currency** field.

2 On the PR Lines tab, type:

Quantity: 25

Unit Cost: 1.23 (Euros)

Using the active exchange rates of 1.23059 USD and 1.5695 CND for the EUR, the system calculates and displays these values on the PR tab:

Total Cost: (25 x 1.23 Euros)= **30.75** (Euros)

Total Base Cost: (30.75 Euros x 1.23059 USD)
= **37.84** (USD)

Total Base 2 Cost: (30.75 Euros x 1.5695 CND)
= **48.26** (CND)

Base 2 currencies are optional. If you configure only one base currency in Organizations, no fields appear for a second base currency.

Configuring Multiple Base Currencies

These guidelines ensure that all appropriate currencies and exchange rates are in the database and that all affected applications perform the correct calculations.

1 In the Currency Codes application, create records as needed for the currency codes to use for Bases 1 and 2.

A Base 1 currency is established on installation. You can type a currency record for Base 2. For example, create a currency record for Euros to use for Base 2.

You can use an existing currency code as the Base 2 currency code.

2 In the Exchange Rates application, open the record for the Base 1 currency code. In the exchange table window, type the Base 2 Currency Code in the **Convert To** field.

For example, if French Francs (FF) is the Base 1 currency code and EUR is the Base 2 currency code:

a Type FF in the **Convert From** field.

b Type EUR in the **Convert To** field.

c Type active and expiration dates for EUR.

3 Type rates for each transaction currency to be converted to the Base 2 currency code.

For example, if French Francs (FF) is the Base 1 currency code and EUR is the Base 2 currency code, and the transaction currency is DEM:

a. Retrieve the DEM record (so DEM is in the **Convert From** field).

b. Type EUR in the **Convert To** field.

c. Type valid active and expiration dates for EUR.

The system can now process transactions in two base currencies. You can create financial reports in both currencies.

Configuring the System With Multiple Languages

A

Overview

A single database can contain data in multiple languages, which lets diverse users run the system in their native language. By default, multiple languages are enabled for:

- Data Dictionary tables
- Company and Item objects
- System messages

For information about running reports in multiple languages, see the *Report Developer Guide*.

The MAXATTRIBUTE table tells you which tables and columns are:

- 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

- Multiple language enabled (MLINUSE=1)

  Indicates whether the column is enabled to store the values in multiple languages

To view tables and columns enabled for multiple language, open a SQL editor and type:

```
select objectname,attributename from maxattribute where mlinuse= 1;
```

These tables and columns are multiple language enabled by default.

Multiple Language Tables

TABLES	COLUMNS
ALNDOMAIN	DESCRIPTION
ASSETATTRIBUTE	DESCRIPTION
COMMTEMPLATE	MESSAGE
COMMTEMPLATE	SUBJECT
COMPANIES	NAME
COMPANIES	NAME_LONGDESCRIPTION

TABLES	COLUMNS
CTRLCONDPROP	PROPERTYVALUE
ITEM	DESCRIPTION_ LONGDESCRIPTION
ITEM	DESCRIPTION
MAXAPPS	DESCRIPTION
MAXATTRIBUTE	TITLE
MAXATTRIBUTE	REMARKS
MAXATTRIBUTECFG	TITLE
MAXATTRIBUTECFG	REMARKS
MAXDOMAIN	DESCRIPTION
MAXLABELS	VALUE
MAXMENU	HEADERDESCRIPTION
MAXMESSAGES	EXPLANATION
MAXMESSAGES	ADMINRESPONSE
MAXMESSAGES	BUTTONTEXT
MAXMESSAGES	SYSTEMACTION
MAXMESSAGES	VALUE
MAXMESSAGES	OPERATORRESPONSE
MAXMODULES	DESCRIPTION
MAXOBJECT	DESCRIPTION
MAXOBJECTCFG	DESCRIPTION
MAXSERVICE	DESCRIPTION
NUMERICDOMAIN	DESCRIPTION
PALETTEITEM	DESCRIPTION
REPORT	DESCRIPTION
REPORTLABEL	LABELVALUE
REPORTLABEL	FONTNAME
REPORTLOOKUP	LABELOVERRIDE
SIGOPTION	DESCRIPTION
SOLUTION	DESCRIPTION
SYNONYMDOMAIN	DESCRIPTION

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

Enabling Multiple Languages on Objects and Attributes

You can enable multiple languages on objects or attributes:

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.

Creating Language Objects

This procedure creates the language object:

1. In the Database Configuration application, select the attribute 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.

NOTE 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 like WO, PO, PR, RFQ, and INVOICE do not.

4. Configure the database. For more details, see *Database Configuration* on 65.

The language tables are empty until you populate them with data. The system provides a toolset to export and import all translatable strings through XLIFF files. For more information, see *Multiple Language Utilities* on 264.

Displaying Non-English Characters

Install additional language files if you find that foreign language characters do not display consistently in the system user interface.

Install the files only if you need them because they require hard disk space and can slow performance when you type text.

Additional Language Options

Option	Languages files installed	Disk space required
Install files for East Asian languages	Chinese, Japanese, and Korean	230 MB
Install files for complex script and right-to-left languages (including Thai)	Arabic, Armenian, Georgian, Hebrew, the Indic languages, Thai, and Vietnamese	10 MB

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

1. From the Start menu, choose **Settings > Control Panel > Regional and Language Options.**
2. Click the **Languages** tab.
3. Select one of the following check boxes to choose the files that you want to install:
   - **Install files for complex script and right-to-left languages (including Thai)**
   - **Install files for East Asian languages**
4. Click **OK** or **Apply**.
5. Restart the computer.

Multiple Language Utilities

The Translation Data Toolkit (TD Toolkit) provides you with the option to use non-English-language databases. The TD Toolkit is a utility that you run to extract translatable data from your database. You can apply 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 add the data.

The following languages are supported:

- Chinese (Simplified)
- Chinese (Traditional)
- Czech
- Dutch
- French
- German
- Hungarian
- Italian
- Japanese
- Korean
- Polish
- Portuguese
- Russian
- Spanish
- Swedish

This document references XLIFF files. XLIFF is an XML file with specific tags for the translation process. For more information, go to the Oasis Web site and search for XLIFF specifications:

http://www.oasis-open.org/home/index.php

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

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

Localizing the Database for an Unsupported Language

If the database is intended for use mainly 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 use mainly by Arabic-speaking users. You can also add a secondary language, such as French, to the new Arabic language database. You can use the TD toolkit to localize the database for an unsupported language.

To localize the database for an unsupported language:

1. Export translatable data from the database.

   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. Translate all XLIFF files from English to the localized language.

   a. Open each XLIFF file and find each instance of the `<target> </target>` tags.

   b. Replace the text between the tags with the text in the language to which you are translating. The text that you enter must match exactly each time you translate a particular word.

3. Import the translated XLIFF files into the database.

   a. Create the following subdirectory:

   `\tools\maximo\lc\xliff`

   `lc` is the two-letter language code.

   You can look up the two-letter language code in the LANGUAGE table, MAXLANGCODE column.

   b. Go to the \tools\maximo\xliff\export directory and copy the files from this directory to the new subdirectory.

c In the installation directory, run the following import command to translate the XLIFF file into the database:

```
tdtoolkit -import -tllc -versionV7100-000
```

`lc` is the two-letter language code.

For example, to import a Hebrew (unsupported) database, run the following import command:

```
tdtoolkit -import -tlhe -version7100-000
```

Adding Second Languages to the Database

You can use the Translation Data toolkit to add non-English language databases. The base language remains the same, but the new language is added as a second language. For example, a Canadian company can add French as a second database language. Use this option to allow non-English language users to view and make database changes. You can add supported and unsupported languages.

Adding Unsupported Second Languages to the Database

To add an unsupported second language to the database:

1 Export translatable data from the database.

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 Translate all XLIFF files from English to the localized language.

a Open each XLIFF file and find each instance of the `<target> </target>` tags.

b Replace the text between the tags with the text in the language to which you are translating. The text that you enter must match exactly each time that you translate a particular word.

3 Add the translated XLIFF files into the database.

a Create the following subdirectory:

```
\tools\maximo\lc\xliff
```

`lc` is the two-letter language code.

You can look up the two-letter language code in the LANGUAGE table, MAXLANGCODE column.

b Go to the \tools\maximo\xliff\export directory and copy the files from the installation directory to the new subdirectory.

c Go to the installation directory and run the following command to translate the XLIFF file into the database:

```
tdtoolkit -addlanglc
```

`lc` is the two-letter language code.

For example, to add Hebrew (an unsupported language) as a secondary language, run the following command:

```
tdtoolkit -addlanghe
```

Adding Supported Second Languages to the Database

When the product was shipped to you, IBM supplied you with all the files that you need to add supported secondary languages to your database.

1 Create the following subdirectory:

```
\tools\maximo\lc\xliff
```

`lc` is the two-letter language code

You can look up the two-letter language code in the LANGUAGE table, MAXLANGCODE column.

2 Go to the \tools\maximo\xliff\export directory and copy the files from the installation directory to the new subdirectory.

3 Add the files that were shipped with the product to the new subdirectory.

4 Go to the installation directory and run the following command to translate the XLIFF file into the database:

```
tdtoolkit -addlanglc
```

`lc` is the two-letter language code.

For example, to add French (a supported language) as a secondary language, run the following command:

```
tdtoolkit -addlangfr
```

Running deletelang.bat

Run this utility to delete a language from the database. You cannot delete the base language.

All table data is removed from the corresponding multiple language tables.

Example

```
deletelang.bat -l FR
```

Running resetbaselang.bat

Run this utility to switch base languages. Example: from English (EN) to French (FR).

The language code is passed as a parameter, and resetbaselang.bat uses the import file language strings to populate the base language strings. This file deletes the current base language and imports a new language from the XML file.

To make the current base language a secondary language:

1. Export the base language.
2. Run the resetbaselang.bat utility.
3. Reimport the original base language.

Example

```
resetbaselang.bat -infile c:\temp\fr.xml -l FR
```

Tracking and Translating New Records in the Base Language

If your system has multiple language implementations, track and perform translations on new records. No auto-translation takes place and therefore, by default, the system stores your new records in the base language only.

To translate your records to the secondary language, choose one of the following two options:

- Translate each record in the localized application.

  Individual record translations should take place only on system implementations that require only a small number of translations.

- Translate your records through the resulting XLIFF file from the TDToolkit.bat utility. For instructions, see *Localizing the Database for an Unsupported Language* on page 265.

- Translate you records through the resulting XLIFF file from the TDToolkit.bat utility. See *Adding Second Languages to the Database* on page 266 for additional information. Because you already translated the XLIFF files, begin at step 3, Add the translated XLIFF files into the database.

- Translate your records through the resultant XML file from the exportlang.bat utility.

Translating Records through the Application

To translate your records to the secondary language through the application, complete the following steps:

1 On the login screen, select the appropriate secondary language and log in to the application.

NOTE Once you select the appropriate secondary language and log in, the language is set.

2 Open the respective application that houses the records in question. For example, from the Start Center select **Go To > Inventory > Item Master**.

3 Select each record that you want to edit. For example, in the Item Master application, type in an item number in the **Item** field and press **Enter**.

4 Change the record and click Save Item. The system saves the changes to secondary language ITEM table.

Translating through the XLIFF Files

To translate your records to the secondary language through the XLIFF file, see *Adding Second Languages to the Database* on 266.

Tracking and Translating Customizations in the Base Language

If your system has multiple language implementations, track and perform translations on system table customizations.

Some system tables that you can customize include:

System Table Examples

Table	Description
MAXATTRIBUTE	Stores information 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

No auto-translation takes place and, therefore, by default, the system stores your system table customizations in the base language only. You translate your system customizations through the XML file from the exportlang.bat utility.

- Translate your records through the resulting XLIFF file from the TDToolkit.bat utility. For instructions, see *Localizing the Database for an Unsupported Language* on page 265.

- Translate you records through the resulting XLIFF file from the TDToolkit.bat utility. See *Adding Second Languages to the Database* on page 266 for additional information. Because you already translated the XLIFF files, begin at step 3 (is this right step #). Add the translated XLIFF files into the database.

Translating through the XLIFF Files

To translate your records to the secondary language through the XLIFF file, see *Adding Second Languages to the Database* on 266.

System Properties Listing

B

The System Properties application manages the system properties. However, to connect to the product's configured database, maximo.properties file contains some properties that are required for the product's EAR deployment in an application server.

Maximo.Properties File

The following table lists the properties in the maximo.properties file.

Property Name	Description
mxe.name	The application server binding the application server object to the RMI registry. The default name is MXServer.
mxe.rmi.port	RMI communication port. If set at zero, RMI uses any available port. You can select another available port number.
mxe.db.user (IBM® DB2®)	Database user the server uses to attach to the database server. For IBM DB2, this user must be an O/S user.
mxe.db.user (Oracle®)	Database user the server uses to attach to the database server. This user must be the schema owner. The default value is maximo.
mxe.db.user (SQL Server)	Database user the server uses to attach to the database server. For SQL Server, this user must have a system administrator role as defined through sp_addsrvrolemember. For example, mxe.db.user = MAXIMO.
mxe.db.password	Password for the database user name.
mxe.db.schemaowner (IBM DB2)	Owner of the database schema. For IBM DB2, the default owner name is maximo.
mxe.db.schemaowner (Oracle)	Owner of the database schema. For Oracle, the default owner name is maximo.
mxe.db.schemaowner (SQL Server)	Owner of the database schema. For SQL Server, the ownername must be dbo.

Property Name	Description
mxe.db.url (IBM DB2)	The default URL is: `mxe.db.url=jdbc:db2://localhost:50000/dbalias` Where `dbalias` is the name of your database.
mxe.db.url (Oracle)	The default URL is: `mxe.db.url=jdbc:oracle:thin:@dbserver:1521:sid` where `dbserver` is the server name of your database server, 1521 is your default Oracle port number, and `sid` is your Oracle system identifier.
mxe.db.url (SQL Server)	The default is: server name, port number, database name defined as: `mxe.db.url=jdbc:inetdae7a:servername:1433?database=databasename&language=us_english& nowarnings=true` where you substitute your database server name, and database name for the italicized values, and 1433 is your 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, the system only supports ASCII for SQL Server.
mxe.db.driver (IBM DB2)	The thin driver defined in mxe.db.driver. `mxe.db.driver=com.ibm.db2.jcc.DB2Driver`
mxe.db.driver (Oracle)	The thin driver defined in mxe.db.driver. For example: `mxe.db.driver=oracle.jdbc.driver.OracleDriver`
mxe.db.driver (SQL Server)	The thin driver defined in mxe.db.driver. For SQL Server, the driver name must be: `mxe.db.driver=com.inet.tds.TdsDriver`

Workflow Properties

The following table lists workflow properties.

You can also see the *Workflow Implementation Guide* for additional information.

Property Name	Description
mxe.workflow.admin	E-mail account of the workflow administrator.

Reorder Properties

The following table lists reorder properties.

Property Name	Description
mxe.reorder.previewtimeout	The reorder preview time out period (in minutes), which should be similar to the Web server session time out. The default value is 30 minutes.

Security Properties

The following table lists security properties.

Maxtype CRYPTO identifies attributes that can be encrypted and decrypted. Maxtype CRYPTOX identifies attributes that can be encrypted, but not decrypted. Each of these maxtypes has its own means of encryption, the parameters for which are defined in the properties file.

Parameters identified as mxe.security.crypto… are for the CRYPTO maxtype, and parameters identified as mxe.security.cryptox… are for the CRYPTOX maxtype.

Property Name	Description
mxe.security.provider	The security provider is obtained from the policy file, which is normally `com.sun.crypto.provider.SunJCE`. To use a different provider, you can specify a value for this parameter.
mxe.security.crypto.mode	The following mode components are valid (OFB must use NoPadding): CBC: Cipher Block Chaining Mode, as defined in FIPS PUB 81. CFB: Cipher Feedback Mode, as defined in FIPS PUB 81. ECB: Electronic Codebook Mode, as defined in 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. OFB: Output Feedback Mode, as defined in FIPS PUB 81. PCBC: Propagating Cipher Block Chaining, as defined by Kerberos V4.
mxe.security.crypto.padding	The following padding components are valid: NoPadding: No padding. PKCS5Padding: The padding scheme described in: RSA Laboratories, *PKCS #5: Password-Based Encryption Standard*, version 1.5, November 1993.
mxe.security.crypto.key	Its length must be a multiple of 24.
mxe.security.crypto.spec	Its length must be a multiple of 8.

Property Name	Description
mxe.security.cryptox.mode	The following mode components are valid (OFB must use NoPadding): CBC: Cipher Block Chaining Mode, as defined in FIPS PUB 81. CFB: Cipher Feedback Mode, as defined in FIPS PUB 81. ECB: Electronic Codebook Mode, as defined in 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. OFB: Output Feedback Mode, as defined in FIPS PUB 81. PCBC: Propagating Cipher Block Chaining, as defined by Kerberos V4.
mxe.security.cryptox.padding	The following padding components are valid: NoPadding: No padding. PKCS5Padding: The padding scheme described in: RSA Laboratories, *PKCS #5: Password-Based Encryption Standard*, version 1.5, November 1993.
mxe.security.cryptox.key	Its length must be a multiple of 24.
mxe.security.cryptox.spec	Its length must be a multiple of 8.
mxe.security.crypto.algorithm	Algorithm is the basic type of encryption used by the system. Crypto properties are used for datatype (maxtype) CRYPTO. The mxe.security.crypto.algorithm property can override the algorithm default value (DESede).
mxe.security.crypto.modulus	Modulus is used only for the RSA algorithm. Crypto properties are used for datatype (maxtype) CRYPTO.
mxe.security.cryptox.algorithm	Algorithm is the basic type of encryption used by the system. Cryptox properties are used for maxtype CRYPTOX (an undecryptable version of crypto). The mxe.security.cryptox.algorithm property can override the algorithm default value (DESede).
mxe.security.cryptox.modulus	Modulus is used only for the RSA algorithm. Cryptox properties are used for maxtype CRYPTOX (an undecryptable version of crypto).
mxe.security.crypto.algorithm	Algorithm is the basic type of encryption used by the system. Crypto properties are used for datatype (maxtype) CRYPTO. The mxe.security.crypto.algorithm property can override the algorithm default value (DESede).
mxe.security.crypto.modulus	Modulus is used only for the RSA algorithm. Crypto properties are used for datatype (maxtype) CRYPTO.

Additional Encryption Algorithms

The default encryption algorithm is DESede. However, some customers require a stronger algorithm for encryption. For these customers, an alternate algorithm can be configured within the additional.properties file. Different properties can be configured for the CRYPTO and CRYPTOX data types.

The following table lists supported encryption algorithms:

Algorithm	Provider	Additional comments
AES	Cryptix, Sun	For Sun® Microsystems, Inc., use mode = ECB.
Blowfish	BouncyCastle, Cryptix	
CAST5	Cryptix	
DES	Cryptix, Sun	
DESede	Cryptix, Sun	
IDEA	Cryptix	
MARS	Cryptix	
PBEWithMD5AndDES	Sun	For Sun Microsystems, Inc., 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	
SKIPJACK	Cryptix	Spec length must be a multiple of 10.
Square	Cryptix	
Twofish	Cryptix	

Debugging Properties

The mbocount, logSQLTimeLimit and the fetchResultLogLimit logging utilities are enabled in the properties file by default. 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

These features are for testing and debugging purposes. When you are satisfied with your system deployment, you can turn off the system performance logging. To disable the logging utilities, modify the properties file to the settings indicated in the following table.

Property Name	Description
mxe.mbocount	Displays the number of business objects created by the server. The default is 1. You can change the value to 0. To disable, edit the file to read `mxe.mbocount=0`.
mxe.db.logSQLTimeLimit	The system logs the SQL statements that take longer than the specified time limit. The time is measured in milliseconds (thousandths of a second). The default value is 1000 milliseconds. To disable, edit the file to read: `mxe.db.logSQLTimeLimit=0`.
mxe.db.fetchResultLogLimit	When this setting 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. The default is 200 rows. To disable, edit the file to read: `mxe.db.fetchResultLogLimit=0`.
mxe.db.logSQLPlan (Oracle only)	Setting this property to `true` logs the execution plan for all SQL statements containing a full table scan. ▼ If you define `mxe.db.sqlTableScanExclude` (as shown in the next example under the Property name column) the system logs all the tables except for the ones you intentionally exclude. ▼ If you do not define `mxe.db.sqlTableScanExclude`, the system logs only the SQL statements that exceed the time limit set in mxe.dblogSQLTimeLimit.

Property Name	Description
mxe.db.sqlTableScanExclude=ACTION,MAXROLE,SCCONFIG,MAXUSER (Oracle only)	You can define the table names which you want to exclude from the log. The table names must be UPPER case. ▼ If you define mxe.db.sqlTableScanExclude, the system logs all the tables except for the ones that you list. ▼ If you do not define `mxe.db.sqlTableScanExclude` – but you do set `mxe.db.logSQLPlan=true` – the system logs only the SQL statements that exceed the time limit set in `mxe.dblogSQLTimeLimit`.

Additional Debugging Parameter

If you decide to use the additional bugging parameter, add it to maximo.properties.

Property Name	Description
mxe.debug.spid=yes	Add this parameter if you want log files to include user names and process ID. This parameter lets you trace SQL statements and blocks to specific users. For IBM DB2 database only: to trace users you need authorization to access the SYSIBM.SYSDUMMY1 table. For Oracle database only: to trace users you need authorization to access the v$session table. Go to <*Maximo root*> applications\Maximo\properties\logging.properties, and ensure `log4j.maximo.sql=INFO`.

BIRT Report Server

The following table lists BIRT report server properties.

Property Name	Description
mxe.report.birt.maxconcurrentrun	Maximum number of reports that can be run concurrently. The default value is 5.
mxe.report.birt.queueidletimeseconds	Number of seconds that the Report Queue Manager is idle after each run. The default value is 60.

Report Integration Properties

The following table lists Business Objects report integration properties

Property Name	Description
mxe.report.bo.db.connnectstring	The Maximo database connection string (Oracle) or ODBC DSN (SQL Server and DB2) defined on the BusinessObjects™ Enterprise server. The default is `mxe.report.bo.db.connectstring=MAXIMO`.
mxe.report.bo.db.databaseName (SQL Server, only)	The database name for the Maximo database. The default is MAXIMO.
mxe.report.bo.serverURL	The URL of the BusinessObjects Enterprise server, including port number and folder. The system uses this URL to access the bocrystal.war Web application. The default URL is http://BOserver:8080/bocrystal.
mxe.report.bo.rootFolder	The BusinessObjects Enterprise root folder name. This property must be rpt, unless you specified another value when you added reports to the central management console. The default is `rpt`.
mxe.report.bo.rptServerLogonName	The BusinessObjects Enterprise logon name. This user must have specific rights to any report you access from the system.
mxe.report.bo.rptServerLogonPass	The BusinessObjects Enterprise password.
mxe.report.bo.cmsName	Unless you changed the name of the central management server when you installed BusinessObjects Enterprise, this property is the name of the server where you installed BusinessObjects Enterprise and the server port number., separated by a colon. To verify the correct values, open the central management server logon page and check the **System** field. For example, BOSERVER:6400.

Customer Report Integration Properties

The following table lists customer report integration properties

Property Name	Description
mxe.report.custom.rptServerLogonPass	Password used to log in to external report server.
mxe.report.custom.serverURL	URL of the custom reporting application.

Depending on your external reporting system, you might need to pass additional property values from the system. To pass additional property values, add the additional properties in the System Properties application.

Other Report Properties

The following table lists other report properties.

Property Name	Description
mxe.report.reportsInAPage	Determines the number of reports that display in the Reporting window that is accessed from system applications. The default value is 5.
mxe.activex	Enables printing of attached documents that are Microsoft® file types (such as .xls, .doc, .ppt). The default value is Y. If users do not want to enable Active X Controls to print Microsoft documents, they must set the value to N (0), which results in no Microsoft documents being used with DPA functionality.

Cron Task Manager Properties

The following table lists cron task manager properties.

Property Name	Description
mxe.crontask.donotrun	Use ALL to exclude all cron tasks from running. To exclude a specific cron task from running, specify the instance by `crontaskname.instancename`.
mxe.cronTaskInitDelay	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. The default value is 60 seconds.
mxe.cronTaskMonitorInterval	The cron task manager monitors the statuses of the cron tasks at the time intervals set by this property. The default value is 60 seconds.
mxe.crontask.historycleanuprate	How often, in minutes, the excessive cron task history records are removed. No action if 0. The default value is 180 minutes.

E-Signature Properties

The following table lists e-signature properties

Property Name	Description
mxe.esig.defaultuserid	Set this flag to true is you want the Esignature login dialog to default to the login ID. The default value is true.

LDAP Integration Properties

The following table lists LDAP integration properties.

Property Name	Description
mxe.allowLDAPUsers	Indicates whether LDAP users are allowed into the system, if they do not have a user record. The default value is 0.
mxe.LDAPGroupMgmt	Indicates whether LDAP owns security group management when mxe.useAppServerSecurity = 1. The default value is 1.
mxe.LDAPUserMgmt	Indicates whether LDAP owns user management when mxe.userAppServerSecurity = 1. The default value is 1.
mxe.ClientCountMinutes	Interval for counting sessions. The default value is 15 minutes.

System Properties

System properties are properties that are managed exclusively using the System Properties application. Most of these properties can be dynamically modified and refreshed without have to shutdown and restart the system's application server. Some of these properties can be modified using the System Properties application, but do not take affect until the next time you restart the application server.

Property Name	Description
mxe.allowLocalObjects	Set to true in production environments, to improve system performance. Set to false for development work, or for custom applications. The default is false.
mxe.useAppServerSecurity	By default you use the system's security, so the value is false. Set to true if you configure the system to use application server provided security.

Property Name	Description
mxe.MLCacheLazyLoad	By default, the multi-language metadata cache loads one object at a time. Set this flag to 1 to load all objects simultaneously for one language.
mxe.UserLicenseKey	The product enabler (license key) is used during installation. If the product enabler changes this value must be updated.
mxe.enableConcurrentCheck	Setting this property to true (1) prevents multiple logins on the same user account. Before you create users, set this property to 1.

Server Properties

The following table lists server properties.

Property Name	Description
mxe.adminuserid	The administrative user. Used by the server for administrative tasks and to run cron tasks. This user must have access to all Sites in the system.
mxe.system.reguser	User registration login name for registering a new user. User name specified must have authorization to create new users. This value is asked for during installation.
mxe.system.regpassword	User registration login password. This value is asked for during installation.
mxe.adminEmail	E-mail address used if the user has not specified an e-mail address in the labor record. This value is asked for during installation.
mail.smtp.host	Name of the host running the SMTP server. This name is needed for facilities that make use of e-mail such as Workflow notifications, Actuate e-mailing, and any error message notifications. Your network administrator can provide this address.
mxe.adminPassword	Password for the administrative user. The product has a default value, maxadmin.
mxe.adminusercredential	The credential of the administrative user.
mxe.adminuserloginid	The system login ID for the administrative user. The default value is maxadmin, the system login ID of the administrative user.
mxe.adminmode.logoutmin	The number of minutes end users have to 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). The default value is 5 minutes.

Property Name	Description
mxe.adminmode.numsessions	The number of administrative sessions allowed once the application server is placed in Admin mode. The default value is 5.
mxe.com.port	Com port.
mxe.convertloginid	Identifies whether login id that the user entered must be converted to uppercase before it is validated. The default value is 0 (do not convert). For conversion, set the value to 1.
mxe.email.charset	The character set for e-mail notifications sent from the product. When this property is defined, it is the charset that is used to encode the subject and message when an e-mail notification is sent. No default value.
mxe.maxsequencecheck	Indicates whether to check for multiple sequence values before using. The default value is 0.
mxe.registry.bindcount	The retry count for the system 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. The default value is 100.
mxe.registry.port	Server RMI Registry port that the server components uses. The default value is 1099.
mxe.usermonitor.timeout	Session timeout period (minutes) on the server to clear the cache. The default value is 30 minutes.
mxe.userrestrictionlrucachesize	Manages the number of entries in the LRU cache that stores the restriction entries. The default value is 1000.

Related Database Properties

The following table lists database-related properties.

Property Name	Description
mxe.db.initialConnections	Number of database connections to create when the application server is started. The default value is 8.
mxe.db.maxFreeConnections	Maximum number of free database connections available in the connection pool. The default value is 8.

Property Name	Description
mxe.db.minFreeConnections	Minimum number of free database connections needed in the connection pool in order for more connections to be allocated. The default value is 5.
mxe.db.newConnectionCount	Number of new connections to be created when the minimum free connections are available in the connection pool. The default value is 3.
mxe.db.transaction_isolation	The system install sets the value to: TRANSACTION_READ_COMMITTED. This value cannot be edited.
mxe.db.format.upper	This value defines the database uppercase function for the system. The default value cannot be edited.
mxe.db.format.date	This value tells the system the database date function. A value of `none` tells the system to pass through the date value. The default value cannot be edited.
mxe.db.format.time	This value tells the system the database time function. A value of `none` tells the system to pass through the time value. The default value cannot be edited.
mxe.db.format.timestamp	This value tells the system the database time stamp function. A value of `none` tells the system to pass through the time stamp value. The default value cannot be edited.
mxe.db.autocommit	This value sets the autocommit mode used for the Write connections. Can be either true or false. The default is false, and the default value cannot be edited.
mxe.db.systemdateformat (IBM DB2)	System date format. For IBM DB2, the value is `current timestamp`.
mxe.db.systemdateformat (Oracle)	System date format. For Oracle, the value is sysdate, and the default value cannot be edited.
mxe.db.systemdateformat (SQL Server)	System date format. For SQL Server, the value is `getdate()`.
mxe.db.format.nullvalue (IBM DB2)	The database-specific format of the nullvalue function. For IBM DB2 the value is COALESCE, and the default value cannot be edited.
mxe.db.format.nullvalue (Oracle)	The database-specific format of the nullvalue function. The value for Oracle is NVL, and the default value cannot be edited.
mxe.db.format.nullvalue (SQL Server)	The database-specific format of the nullvalue function. The value for SQL Server must be set to ISNULL.
mxe.db.sqlserverPrefetchRows (SQL Server only)	Setting to reduce lock contention. Optimal setting is 200 rows. Setting a value larger than 500 can degrade performance. The default value is 0.
mxe.db.disableservercursor	Disable the server cursor. The default value is 1.

Property Name	Description
mxe.db.fetchsize	The size of the database fetch. The default value is 40.
mxe.db.fetchsizeuse	Flag to indicate whether to use the fetchsize. The default value is 1.
mxe.db.optionnum	Size of the option. The default value is 1000.
mxe.db.optionuse	Flag to indicate whether to use the option. The default value is 1.
mxe.db.proxyauthentication.mode	The oracle proxy authentication mode is only valid when you are using Oracle Proxy DataBase Manager. Values include: ▼ 1 = username ▼ 2 = username + password ▼ 3 = distinguished name (DN) ▼ 4 = certificate
mxe.dbmanager	This references the Java™ class of the Maximo database manager. The default value is psdi.server.DBManager. If you have an Oracle database requiring proxy authentication, set this property to psdi.server.OracleProxyDBManager. This property also requires you to: ▼ Specify the jdbc database connection string as the oci connection string ▼ Make the Oracle oci driver accessible to the system Web component JVM
mxe.db.refcount	The reference count for the connection logging. The default value is 100

Property Name	Description

mxe.db.resultsettype	`TYPE_FORWARD_ONLY`
	
	`public static final in TYPE_FORWARD_ONLY`
	
	The constant indicates the type for a ResultSet object whose cursor can only move forward.
	
	`TYPE_SCROLL_INSENSITIVE`
	
	`public static final int TYPE_SCROLL_INSENSITIVE`
	
	The constant indicates the type for a ResultSet object that is scrollable, but generally not sensitive to changes that others made.
	
	`TYPE_SCROLL_SENSITIVE`
	
	`public static final int TYPE_SCROLL_SENSITIVE`
	
	The constant indicates the type for a ResultSet object that is scrollable and generally sensitive to changes that others made.
	
	The default value is TYPE_FORWARD_ONLY.

mxe.db.retrydbconnection	Retry getting the database connection when starting the MXServer. The default value is 0.
mxe.db.rowcount	The SQLServer row count value. The default value is 0.

Migration Manager Properties

The following table lists migration manager properties

Property Name	Description
mxe.dm.dmroot	Migration Manager root folder name that is on the application server computer. This folder stores Migration Manager package files. There is not a default value. Users must configure this property before using Migration Manager.
mxe.dm.dmsessiontimeout	Migration Manager http session timeout value. When a long-running Migration Manager task (such as package creation or package deployment) launches, the time-out value for the session of the user who is currently logged in changes to the value specified by this property. The default value is 120 minutes.

Property Name	Description
mxe.dm.dmstagecommit	This value specifies the commit interval for records that are inserted into the Migration Manager's staging table in the target environment. The value is specified in source, and Migration Manager uses the value when distributing a package from the source database to the target database. The default value is 1 (indicating that the commit is performed for every one record).

Attached Documents Properties

The following table lists attached documents properties

Property Name	Description
mxe.doclink.doctypes.defpath	Default path for the doclinks folder on the application server computer. Where the physical documents that are attached to a system record are stored. No default value. An administrator must configure before using the system. Only needs a Live Refresh.
mxe.doclink.maxfilesize	Maximum file size (MB) for doclinks that can be uploaded. The default value is 10 MB. Can be changed in the System Properties application; however, the EAR file must be rebuilt.
mxe.doclink.multilang.aix.websphere	Whether the system is running on an AIX WebSphere platform. The default value is false. Change the value to true if the system is running on an AIX WebSphere 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.
mxe.doclink.path1 through mxe.doclink.path10	Specify the http server path to link documents that are attached to records. No default value. Only needs a Live Refresh.

The system is delivered with ten properties for doclink path translations; however, you can specify additional properties by adding new properties into the MaxProp table through Property Maintenance. LinkedDocumentInfo loads as many as 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.

TIP The `<PATH>` tag is used in the properties file (because of problems specifying ":"), the ":" character is permitted on the database.

For example:

Current Format	New Property Name
C<PATH>\\Doclinks	mxe.doclink.path1
/home/mxadmin/DOCLINKS	mxe.doclink.path2

The MaxPropValue table for doclinks would contain the values listed in the following table.

Property Name	Property Value
mxe.doclink.maxfilesize	10
mxe.doclink.doctypes.defpath	C:\DOCLINKS\
mxe.doclink.path1	C:\Doclinks=http://documentserver/
mxe.doclink.path2	/home/mxadmin/DOCLINKS=http://documentserver/

Work Order Generation Properties

The following table lists work order generation properties

Property Name	Description
mxe.msgLogFile	Work order generation log file.

Integration Properties

For information about integration properties, see the *Integration Guide*.

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:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, 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 Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites 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.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM
AIX
DB2
developerWorks
Everyplace
ibm.com
Lotus
Maximo
Notes
QuickPlace
Tivoli
WebSphere

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft, Windows, Windows NT, 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.

Index

A

B

C

D

E

F

G

H

I

J

L

M

T

U

V

W

X