



# BEA Tuxedo

## Introducing the BEA Tuxedo System

BEA Tuxedo Release 7.1  
Document Edition 7.1  
May 2000

# Copyright

Copyright © 2000 BEA Systems, Inc. All Rights Reserved.

## Restricted Rights Legend

This software and documentation is subject to and made available only pursuant to the terms of the BEA Systems License Agreement and may be used or copied only in accordance with the terms of that agreement. It is against the law to copy the software except as specifically allowed in the agreement. This document may not, in whole or in part, be copied photocopied, reproduced, translated, or reduced to any electronic medium or machine readable form without prior consent, in writing, from BEA Systems, Inc.

Use, duplication or disclosure by the U.S. Government is subject to restrictions set forth in the BEA Systems License Agreement and in subparagraph (c)(1) of the Commercial Computer Software-Restricted Rights Clause at FAR 52.227-19; subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013, subparagraph (d) of the Commercial Computer Software--Licensing clause at NASA FAR supplement 16-52.227-86; or their equivalent.

Information in this document is subject to change without notice and does not represent a commitment on the part of BEA Systems. THE SOFTWARE AND DOCUMENTATION ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. FURTHER, BEA Systems DOES NOT WARRANT, GUARANTEE, OR MAKE ANY REPRESENTATIONS REGARDING THE USE, OR THE RESULTS OF THE USE, OF THE SOFTWARE OR WRITTEN MATERIAL IN TERMS OF CORRECTNESS, ACCURACY, RELIABILITY, OR OTHERWISE.

## Trademarks or Service Marks

BEA, ObjectBroker, TOP END, and Tuxedo are registered trademarks of BEA Systems, Inc. BEA Builder, BEA Connect, BEA Manager, BEA MessageQ, BEA Jolt, M3, eSolutions, eLink, WebLogic, and WebLogic Enterprise are trademarks of BEA Systems, Inc.

All other company names may be trademarks of the respective companies with which they are associated.

## Introducing the BEA Tuxedo System

---

<b>Document Edition</b>	<b>Date</b>	<b>Software Version</b>
7.1	May 2000	BEA Tuxedo Release 7.1

---

---

# Contents

## 1. BEA Tuxedo System Fundamentals

For More Information.....	1-1
What Is the BEA Tuxedo System.....	1-2
Features of the BEA Tuxedo System .....	1-3
Anatomy of the Client/Server Model .....	1-5
Characteristics of Client/Server Architecture .....	1-5
Differences Between 2-Tier and 3-Tier Client/Server Architectures.....	1-6
Client/Server Variations to Suit Your Needs .....	1-8
How the BEA Tuxedo System Fits into the Client/Server Model.....	1-9
What Are Clients, Servers, and Services in a BEA Tuxedo Environment.....	1-11
What Is a BEA Tuxedo Client.....	1-11
What Is a BEA Tuxedo Server .....	1-12
What Are BEA Tuxedo Services .....	1-12
Services Provided by the BEA Tuxedo System .....	1-12
Administrative Services .....	1-12
Application Processing Services .....	1-13
BEA Family of Products .....	1-14

## 2. BEA Tuxedo System Architecture

Basic Architecture of the BEA Tuxedo System .....	2-1
What You Can Do Using the ATMI.....	2-4
What Are the BEA Tuxedo Messaging Paradigms .....	2-8
What Is Conversational Communication.....	2-9
How the EventBroker Works .....	2-10
What Types of Events Are Reported .....	2-11
How Are Events Reported.....	2-12
What Is Queue-based Communication.....	2-13

---

Using Application Queues.....	2-13
What Is Request/Reply Communication .....	2-14
What Is Synchronous Messaging .....	2-15
What Is Asynchronous Messaging .....	2-16
What Is Unsolicited Communication .....	2-17
What Are Nested and Forwarded Service Requests .....	2-18
Nested Requests.....	2-18
Forwarded Requests .....	2-20
How the BEA Tuxedo System Processes Messages .....	2-21
What Are the Benefits of Service Request Processing.....	2-24
What Are Typed Buffers .....	2-24
Characteristics of Buffer Types.....	2-25
Using the MIB .....	2-29
Types of MIB Users .....	2-30
Classes, Attributes, and States in the MIB .....	2-30
BEA Tuxedo Application Processing Services .....	2-31
What Is Data Compression .....	2-31
What Is Data-dependent Routing .....	2-32
Uses of Data-dependent Routing.....	2-33
Example of Data-dependent Routing with a Horizontally Partitioned Database .....	2-34
Example of Data-dependent Routing with Rule-based Servers .....	2-35
Example of Data-dependent Routing with Distributed Application .....	2-36
What Are Encoding and Decoding of Data .....	2-37
What Is Data Encryption .....	2-38
What Is Data Marshalling .....	2-38
What Is Load Balancing .....	2-39
Assigning a Load Factor.....	2-40
What Is Message Prioritization.....	2-41
What Is Meant by Naming.....	2-42
Naming Services.....	2-42
Advertising Services.....	2-43
Naming Events .....	2-44
BEA Tuxedo Administrative Services .....	2-44

---

### 3. Three Ways of Viewing the BEA Tuxedo System Infrastructure

Basic BEA Tuxedo System Infrastructure .....	3-1
Management View: Using Administrative Tools.....	3-2
Available BEA Tuxedo System MIBs .....	3-3
Using the BEA Administration Console .....	3-4
Browser Requirements .....	3-4
Benefits of Using the BEA Administration Console.....	3-5
Exploring the Main Menu of the BEA Administration Console.....	3-6
Using the Configuration Tool .....	3-7
What Is the Tree .....	3-8
Using the Power Bar .....	3-9
Managing Operations Using the MIB .....	3-10
Types of MIB Users .....	3-11
Classes, Attributes, and States in the MIB .....	3-11
Using Command-Line Utilities .....	3-12
Configuring Your Application Using Command-line Utilities.....	3-12
Operating Your Application Using Command-line Utilities.....	3-14
Managing System Events Using EventBroker .....	3-14
What Is an Event .....	3-15
Subscribing to an Event.....	3-16
Types of Events .....	3-17
Differences between System and Application-specific Events.....	3-18
BEA Tuxedo Administrative Services .....	3-19
Managing Application Queues.....	3-20
Using qmadmin to Administer Application Queues .....	3-20
Using tmconfig to Modify Your Configuration .....	3-21
Managing Your Configuration .....	3-22
Creating the Configuration File.....	3-22
Making Permanent Configuration Changes .....	3-24
Managing Your Configuration Dynamically.....	3-25
Performing Dynamic Operations Using tadmin(1).....	3-26
Commonly Used tadmin Commands .....	3-26
Sample Output from the tadmin Command .....	3-27

---

Managing a Distributed Application Centrally .....	3-28
Managing Security.....	3-30
Selecting Security Options .....	3-31
Setting Up Security.....	3-32
Starting Up and Shutting Down your Application .....	3-33
Managing Transactions.....	3-34
Coordinating Operations with a Transaction Manager Server (TMS) .....	3-35
Tracking Participants with a Transaction Log (TLOG) .....	3-35
Managing Workstations.....	3-36
Development View: What You Can Do Using the ATMI .....	3-37
Run-Time System View: Using Tools in Different Configurations.....	3-41
Run-Time System Capabilities.....	3-42
What Is a Single-machine Configuration .....	3-43
What Is a Multiple-machine (Distributed) Configuration .....	3-45
What Is a Multiple-domain Configuration .....	3-49
Features of a Multiple-domain Configuration .....	3-53
What Is a BEA Tuxedo BRIDGE.....	3-53
What Is the Role of the Bulletin Board and Bulletin Board Liaison .....	3-55
What Are Clients and Servers .....	3-56
What Is the Distinguished Bulletin Board Liaison (DBBL).....	3-57
What Are the Domains Administrative Tools .....	3-57
What Are IPC Message Queues .....	3-59
When to Use Single Server, Single Queues (SSSQ) .....	3-60
When to Use Multiple Server, Single Queue (MSSQ) Sets .....	3-60
What Are the Workstation Handler and Workstation Listener .....	3-62
How a Workstation Client Connects to an Application .....	3-63
What Is the User Log (ULOG).....	3-64
How Is the ULOG Created.....	3-64
Example of a ULOG Message.....	3-64
Where the ULOG Resides .....	3-66

## **4. Integrating the BEA Tuxedo Product Family in an Enterprise System**

BEA Product Integration .....	4-2
BEA Product Suite .....	4-2

---

Mainframe Connectivity: Using BEA eLink.....	4-4
Components of the BEA eLink Product Suite.....	4-5
BEA eLink for Mainframe - TCP/IP for MVS (for IMS and CICS) .....	4-6
BEA eLink for Mainframe - SNA.....	4-7
BEA eLink for Mainframe - OSI TP.....	4-8
Internet Access: Using BEA Jolt.....	4-9
Components of BEA Jolt.....	4-10
Developing and Managing Applications: Using BEA Manager .....	4-11
BEA Manager Components .....	4-12
Components of the BEA Tuxedo Product .....	4-12
Online Transaction Processing: Using the Core BEA Tuxedo System.....	4-13
Enabling Scalability: Using BEA Tuxedo Domains .....	4-15
Features of BEA Tuxedo Domains .....	4-17
What Is a Domain .....	4-18
What Is a Domains Gateway .....	4-18
What Are Domain Gateway Types.....	4-19
BEA Tuxedo Domains Components .....	4-20
Storing Messages and Service Requests: Using BEA Tuxedo /Q.....	4-21
Using the Message Queuing Server .....	4-22
Storing and Forwarding Messages .....	4-23
BEA Tuxedo /Q Capabilities.....	4-25
Workstation Connectivity: Using BEA Tuxedo Workstation .....	4-26
Workstation Components .....	4-27
Developing Client-Server Architecture Using WebLogic Enterprise .....	4-28
Developing and Managing Distributed Java-based Applications: Using BEA	
WebLogic Server.....	4-30
WebLogic Server Implementations.....	4-31
Advantages of Using BEA WebLogic Server.....	4-31





# 1 BEA Tuxedo System Fundamentals

- What Is the BEA Tuxedo System
- Anatomy of the Client/Server Model
- How the BEA Tuxedo System Fits into the Client/Server Model
- What Are Clients, Servers, and Services in a BEA Tuxedo Environment
- Services Provided by the BEA Tuxedo System
- BEA Family of Products

## For More Information

Many resources are available to help you understand the BEA Tuxedo system. The following books, white papers, and presentations provide information about client/server architecture, building and managing distributed business applications, and using the BEA Tuxedo system to build and manage enterprise applications:

- Andrade, Juan, M. Carges, T. Dwyer, and S. Felts, *The Tuxedo System - Software for Constructing and Managing Distributed Business Applications*. Reading, Massachusetts: Addison-Wesley Publishing, 1996.
- Edwards, Jeri, with D. DeVoe, *3-Tier Client/Server at Work*. New York: John Wiley & Sons, Inc., April 1997.
- Edwards, Jeri, D. Harkey, R. Orfali, *The Essential Client/Server Survival Guide*. New York: John Wiley & Sons, Inc., May 1997.

- Hall, Carl, *Building Client/Server Applications Using Tuxedo - Designing and Building Cost-Effective, High Performance Client/Server Applications Using Tuxedo*. Wiley Computer Publishing.
- Lee, Rich, *BEA Tuxedo Essentials*. Presented at the BEA User's Conference in New Orleans, La., February, 1999.
- MacBlane, Randy, *Managing your BEA Tuxedo Applications Even Over the Internet*. Presented at the BEA User's Conference in San Jose, Ca., May 1997.
- MacBlane, Randy, *Tuxedo's Management Information Base*. Presented at the BEA User's Conference in San Francisco, Ca., February 1996.
- *BEA Tuxedo: The Programming Model* (White Paper)
- *BEA Tuxedo and the Component Software Model* (White Paper)
- *Inter-Application Transaction Processing with BEA Tuxedo Domains* (White Paper)
- *Reliable Queuing Using BEA Tuxedo* (White Paper)

## What Is the BEA Tuxedo System

The BEA Tuxedo system is a *middleware* product that distributes applications across multiple platforms, databases, and operating systems using message-based communications and, if desired, distributed transaction processing.

Middleware is used with client/server applications to distribute processing among multiple servers, manage distributed transactions, and integrate multiple database platforms. Middleware systems are sometimes known as “on-line transaction processing” or “OLTP” systems.

The BEA Tuxedo system is a mature product based on over 15 years of development from a diverse group of technology companies including AT&T, UNIX System Laboratories (USL), Novell, and BEA Systems, Inc. It is both a development platform and an execution platform. The BEA Tuxedo system serves as an extension to the operating system.

The BEA Tuxedo system provides the following:

- An industry standard for the creation and central administration of distributed on-line transaction applications in a heterogeneous client/server environment.
- Ease of use for application developers, who do not need to know all the details about server locations, routing, or platforms used. In a BEA Tuxedo application, these aspects of a program are transparent.
- The fundamental underpinnings for creating, managing, and maintaining reliable, high performance, easily managed distributed systems.

## Features of the BEA Tuxedo System

The BEA Tuxedo system offers many features to accommodate the needs of the administrator, architect, and programmer of an application.

### Administrative Features

- *Password security and access control security*—Password security allows application designers to control access by requiring passwords at initialization time (authentication). Further control is available through authorization, a means of restricting access to certain application services to clients that have been given explicit permission and that have authenticated identities.
- *System events notification*—The BEA Tuxedo system provides details about system events, such as servers dying and network failures. When an event is posted by clients or servers, the EventBroker looks up all the subscribers to that event and takes appropriate actions, as determined by each subscription.
- *The MIB (Management Information Base)*—An administrative interface that enables you to monitor, configure, and tune your application through your own programs. It is an implementation-independent management database defined as a set of FML attributes, which allows you to query or change information.
- *Web-based administration*—A graphical user interface, available through the World Wide Web, for the configuration and control of BEA Tuxedo applications.

## Architectural Features

- *Distributed services*—Allow transparent access to application and/or system services located on different hardware platforms.
- *Fast, connectionless communications*—Clients connect to a bulletin board rather than to servers, thus improving system performance.
- *Scalability*—You can quickly scale your application to match varying system load demands because services and servers can be replicated and distributed easily. You can set thresholds programmatically to enable the BEA Tuxedo system to spawn new servers or to shut down servers automatically.
- *Server transparency*—The directory of services on the bulletin board maps service names to servers; clients do not need to be aware of server identity.

## Programming Features

- *Communication techniques*—The Application Programming Interface (API) for the BEA Tuxedo system is a superset of X/Open's XATMI interface called the Application to Transaction Monitor Interface or ATMI. The BEA Tuxedo ATMI is a rich set of communication techniques for writing distributed applications.
- *Distributed Transaction Processing (DTP)*—Allows work being done throughout a distributed application to be atomically completed—an essential characteristic of any OLTP system.
- *Typed buffers*—Provide transparent handling of application data across heterogeneous platforms.
- *X/Open TX compliance*—The BEA Tuxedo system conforms to the X/Open interface standard for transaction demarcation.
- *X/Open XA compliance*—The BEA Tuxedo system conforms to the X/Open interface standard for transaction database systems (called resource managers). As a result, you can mix and match databases within one application while maintaining data integrity.

## See Also

- “Services Provided by the BEA Tuxedo System” on page 1-12
- “BEA Family of Products” on page 1-14

# Anatomy of the Client/Server Model

In client/server architecture, clients, or programs that represent users who need services, and servers, or programs that provide services, are separate logical objects that communicate over a network to perform tasks together. A client makes a request for a service and receives a reply to that request; a server receives and processes a request, and sends back the required response.

## Characteristics of Client/Server Architecture

- *Asymmetrical protocols*—There is a many-to-one relationship between clients and a server. Clients always initiate a dialog by requesting a service. Servers wait passively for requests from clients.
- *Encapsulation of services*—The server is a specialist: when given a message requesting a service, it determines how to get the job done. Servers can be upgraded without affecting clients as long as the published message interface used by both is unchanged.
- *Integrity*—The code and data for a server are centrally maintained, which results in cheaper maintenance and the protection of shared data integrity. At the same time, clients remain personal and independent.
- *Location transparency*—The server is a process that can reside on the same machine as a client or on a different machine across a network. Client/server software usually hides the location of a server from clients by redirecting service requests. A program can be a client, a server, or both.
- *Message-based exchanges*—Clients and servers are loosely-coupled processes that can exchange service requests and replies using messages.
- *Modular, extensible design*—The modular design of a client/server application enables that application to be fault-tolerant. In a fault-tolerant system, failures may occur without causing a shutdown of the entire application. In a fault-tolerant client/server application, one or more servers may fail without stopping the whole system as long as the services offered on the failed servers are available on servers that are still active. Another advantage of modularity is that a client/server application can respond automatically to increasing or decreasing system loads by adding or shutting down one or more services or servers.

- *Platform independence*—The ideal client/server software is independent of hardware or operating system platforms, allowing you to mix client and server platforms. Clients and servers can be deployed on different hardware using different operating systems, optimizing the type of work each performs.
- *Reusable code*—Service programs can be used on multiple servers.
- *Scalability*—Client/server systems can be scaled horizontally or vertically. Horizontal scaling means adding or removing client workstations with only a slight performance impact. Vertical scaling means migrating to a larger and faster server machine or adding server machines.
- *Separation of Client/Server Functionality*—Client/server is a relationship between processes running on the same or separate machines. A server process is a provider of services. A client is a consumer of services. Client/server provides a clean separation of functions.
- *Shared resources*—One server can provide services for many clients at the same time, and regulate their access to shared resources.

## Differences Between 2-Tier and 3-Tier Client/Server Architectures

Every client/server application contains three functional units:

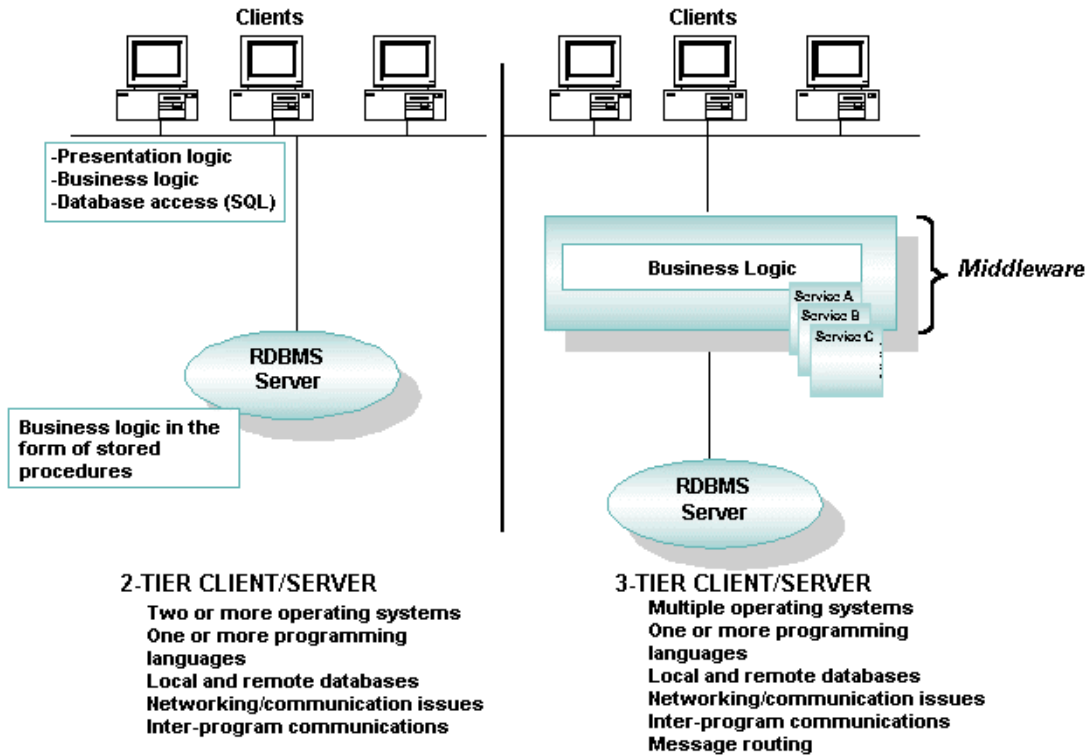
- Presentation logic or user interface (for example, ATM machines)
- Business logic (for example software that enables a customer to request an account balance)
- Data (for example, records of customer accounts)

These functional units can reside on either the client or on one or more servers in your application. Which of the many possible variations you choose depends on how you split the application and which middleware you use to communicate between the tiers.

In 2-tier client/server applications, the business logic is buried inside the user interface on the client or within the database on the server in the form of stored procedures. Alternatively, the business logic can be divided between the client and server. File servers and database servers with stored procedures are examples of 2-tier architecture.

In 3-tier client/server applications, the business logic resides in the middle tier, separate from the data and user interface. In this way, processes can be managed and deployed separately from the user interface and the database. Also, 3-tier systems can integrate data from multiple sources.

Figure 1-1 2-Tier and 3-Tier Client/Server Models



## Client/Server Variations to Suit Your Needs

Client/server architecture can accommodate the needs of each of the following situations:

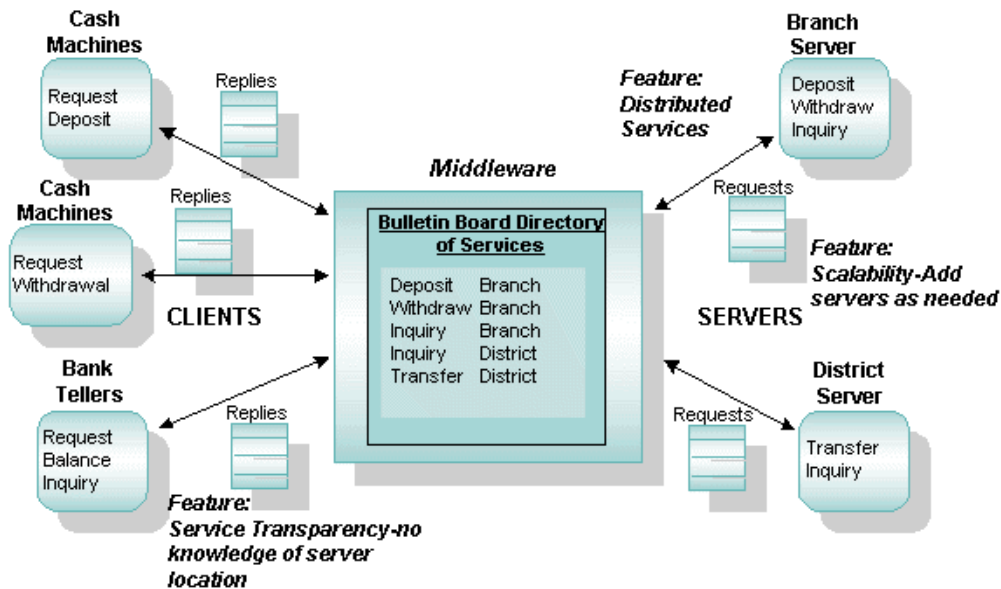
- Small shops and laptops—The client, the middleware software, and most of the business services operate on the same machine. We recommend this approach for one-person businesses such as a dentist's office, a home office, and a business traveler who frequently works on a laptop computer.
- Small businesses and corporate departments—A LAN-based single-server application is required. Users of this type of application include small businesses, such as a medical practice with several doctors, a multi-department corporation, or a bank with several branch offices. In this type of application, multiple clients talk to a local server. Administration is simple: security is implemented at the machine level and failures are detected easily.
- Large enterprises—Multiple servers that offer diverse functionality are required. Multiple servers can reside on the Internet, intranets, and corporate networks, all of which are highly scalable. Servers can be partitioned by function, resources, or databases, and can be replicated for increased fault tolerance or enhanced performance. This model provides a great amount of power and flexibility. How well you architect your application is critical to this client/server model. You may need to partition work among servers, or design servers to delegate work to other servers.



# How the BEA Tuxedo System Fits into the Client/Server Model

The BEA Tuxedo system fits into the middle of the client/server model. In a BEA Tuxedo application, clients log in and request services offered by an application. The BEA Tuxedo system offers these services through a transparent bulletin board. The bulletin board contains a directory advertising services. In a banking application, for example, the bulletin board might advertise deposit, withdrawal, and inquiry services. The BEA Tuxedo system then finds a server (for example, at the appropriate branch or district office) that can provide the requested services.

Figure 1-2 Clients and Servers in a Sample Banking Application



The preceding figure shows the primary building blocks of a BEA Tuxedo application:

- **Clients**—programs that collect input from users, sends requests through the BEA Tuxedo system to servers, and then collects the replies from servers and delivers them to the users.

- *Servers*—programs that encapsulate the business logic into a set of services that define the application.
- *Middleware*—comprises all the distributed software needed to support interactions between clients and servers. It is the medium that enables a client to obtain a service from a server. Middleware includes: API functions used by the client (to issue requests and receive replies) and the server (to issue replies) and messaging paradigms used to transmit client requests and server responses over a network. Middleware does not include any of the following: the user interface on the client, application logic, and services provided by servers.

In this sample BEA Tuxedo application at a bank, clients (cash machines and tellers) make requests, and servers (at branch and district offices) provide services and responses. For example, a customer may use a cash machine to find out how much money is available in his personal checking account. The cash machine (a client) calls the server to get the balance. The server receives the request, retrieves the balance, and sends the information to the cash machine.

## See Also

- “Anatomy of the Client/Server Model” on page 1-5
- “What Are Clients, Servers, and Services in a BEA Tuxedo Environment” on page 1-11

# **What Are Clients, Servers, and Services in a BEA Tuxedo Environment**

This topic describes a client, server, and services in a BEA Tuxedo environment.

## **What Is a BEA Tuxedo Client**

A client is a program that collects a request from a user and passes that request to a server capable of fulfilling it. It can reside on a PC or workstation as part of the front end of an application. It can also be embedded in software that reads a communication device such as an ATM machine from which data is collected and formatted before being processed by BEA Tuxedo servers.

To be a client, a program must be able to invoke the BEA Tuxedo libraries of functions and procedures known collectively as the Application to Transaction Monitor Interface, or ATMI. The ATMI is supported in several language bindings.

A client joins a BEA Tuxedo application by calling the ATMI client initialization routine. Once it has joined an application, a client can define transaction boundaries and call ATMI functions that enable it to communicate with other programs in your application. The client leaves the BEA Tuxedo application by issuing an ATMI termination function. By joining an application only when necessary and leaving it once the appropriate task is complete, a client frees BEA Tuxedo system resources for use by other clients and servers.

When building a distributed application, you must determine how information is gathered and presented to your business for processing. You have complete control over where and when to call ATMI functions, depending upon your business logic and rules. Your program can join one BEA Tuxedo application, perform some tasks and leave, and then join a different BEA Tuxedo application to perform another task. If you are using a multicontexted application, your client can perform tasks in more than one application without leaving any of them.

## What Is a BEA Tuxedo Server

A BEA Tuxedo server is a process that oversees a set of services, dispatching them automatically for clients that request them. A service, in turn, is a function within a server program that performs a particular task needed by a business. A bank, for example, might have one service that accepts deposits and another that reports account balances. A server at this bank might receive requests from clients for both services. It is the server's job to dispatch each request to the appropriate service.

Service functions implement business logic through calls to database interfaces such as SQL and, possibly, calls to the ATMI to access additional services, queues, and other resources. The servers on which these services reside then reply to the clients or forward client requests to a new service.

## What Are BEA Tuxedo Services

A service is a module of application code that performs a task. Services are compiled and link edited to form executable servers.

## Services Provided by the BEA Tuxedo System

The BEA Tuxedo system offers many services to help you streamline and administer your application.

### Administrative Services

The BEA Tuxedo system provides services for the following administrative tasks:

- [Application queue management](#)
- [Centralized application configuration](#)
- [Distributed application management](#)

- Dynamic application reconfiguration
- Event management
- Security management
- Startup and shutdown of an application
- Transaction management
- Workstation management

## **Application Processing Services**

The BEA Tuxedo system provides services that enable you to implement the following functionality in your application:

- Data compression
- Data-dependent routing
- Data encoding
- Data encryption
- Data marshalling
- Load balancing
- Message prioritization
- Service and event naming

# BEA Family of Products

BEA offers the following family of products.

This Product	Provides
<a href="#">BEA eLink</a>	A suite of connectivity products that allow you to seamlessly integrate BEA Tuxedo distributed applications with enterprise applications
<a href="#">BEA Jolt</a>	A BEA Tuxedo client API in Java
<a href="#">BEA Manager</a>	A package of administrative tools for a BEA Tuxedo or WebLogic Server application
The BEA Tuxedo system consists of four components: <ul style="list-style-type: none"><li>■ <a href="#">Core BEA Tuxedo system</a></li><li>■ <a href="#">Domains</a></li><li>■ <a href="#">/Q</a></li><li>■ <a href="#">Workstation</a></li></ul>	<ul style="list-style-type: none"><li>■ <i>The core BEA Tuxedo product</i>—enables you to build high-performance, mission-critical, and reliable distributed applications. It provides the framework for building scalable 3-tier client-server applications in heterogeneous, distributed environments.</li><li>■ <i>Domains</i>—extends the BEA Tuxedo client/server model to provide transaction interoperability across separately administered BEA Tuxedo applications.</li><li>■ <i>/Q</i>—allows reliable queuing of requests.</li><li>■ <i>Workstation</i>—offers full client support for a wide variety of operating systems allowing applications to use remote clients that do not need a full BEA Tuxedo implementation.</li></ul>
<a href="#">BEA WebLogic Enterprise</a>	A leading application server family of products providing businesses and organizations that depend on mission-critical applications with the advantages of Common Object Request Broker Architecture (CORBA)-compliant and Enterprise Java Beans (EJB) programming models, combined with the power, robustness, and proven reliability of the BEA Tuxedo system.
<a href="#">BEA WebLogic Server</a>	A Java-application server for developing, integrating, deploying, and managing large-scale, distributed Web, network, and database applications

# 2 BEA Tuxedo System Architecture

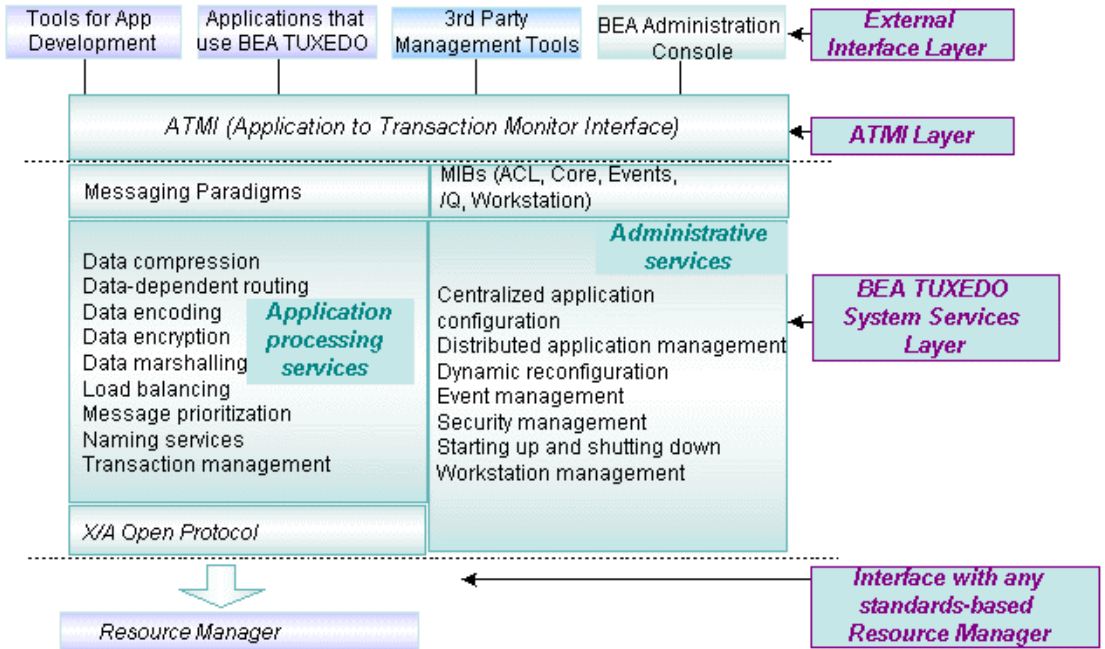
- Basic Architecture of the BEA Tuxedo System
- What Are the BEA Tuxedo Messaging Paradigms
- How the BEA Tuxedo System Processes Messages
- BEA Tuxedo Application Processing Services
- BEA Tuxedo Administrative Services

## Basic Architecture of the BEA Tuxedo System

The following figure illustrates the basic architectural elements of a BEA Tuxedo system: external interfaces to the system, the ATMI layer, the MIB, BEA Tuxedo system services, and the system's interface with standards-compliant resource managers.

## 2 BEA Tuxedo System Architecture

**Figure 2-1 The BEA Tuxedo System Basic Architecture**



As shown in this illustration, the BEA Tuxedo system contains the following parts.

Architectural Part	Description
External interface layer	This layer consists of interfaces between the user and the system. It includes both tools for application development, such as Simple Network Management Protocol (SNMP) agents, and tools for administration, such as the BEA Administration Console. The BEA Administration Console and SNMP agents can interact with standard management consoles. Thus a user can manage a BEA Tuxedo system and a network configuration from one console. In addition, application architects and developers can build their own administrative tools or application- or market-specific tools on top of the MIB.



Architectural Part	Description
<a href="#">ATMI (Application to Transaction Monitor Interface)</a>	The interface between an application and the BEA Tuxedo system. The ATMI and the BEA Tuxedo system implement the X/Open DTP model of transaction processing. An abstract environment, the ATMI supports location transparency and hides implementation details. As a result, programmers are free to configure and deploy BEA Tuxedo applications to multiple platforms without modifying the application code.
<a href="#">Messaging paradigms</a>	Different models of transferring messages between a client and a server. Examples include request/response mode, conversational mode, events and unsolicited notification.
<a href="#">MIB (Management Information Base)</a>	The Management Information Base (MIB) is an interface that enables users to program and administer a BEA Tuxedo system easily. MIB operations enable you to perform all management tasks (monitoring, configuring, tuning, and so on). The MIB allows you to perform one task to one object at a time or to build tool kits with which you can batch tasks and/or objects. (For information about available MIBs, see “Available BEA Tuxedo System MIBs” on page 3-3.)
<a href="#">BEA Tuxedo Services (administrative services and application processing services)</a>	Services and/or capabilities provided by the BEA Tuxedo system infrastructure for developing and administering applications. The application processing services available to developers include: data compression, data-dependent routing, data encoding, load balancing, and transaction management. The administrative services include: centralized application configuration, distributed application management, domains partitioning, dynamic reconfiguration, event and fault management, IPC message queues, and workstation management. (For information on administrative services, see the topic titled: “Three Ways of Viewing the BEA Tuxedo System Infrastructure” on page 3-1.)
<a href="#">Resource Manager</a>	A software product in which data is stored and available for retrieval through application-based queries. The resource manager (RM) interacts with the BEA Tuxedo system and implements the XA standard interfaces. The most common example of a resource manager is a database. Resource managers provide transaction capabilities and permanence of actions; they are the entities accessed and controlled within a global transaction.

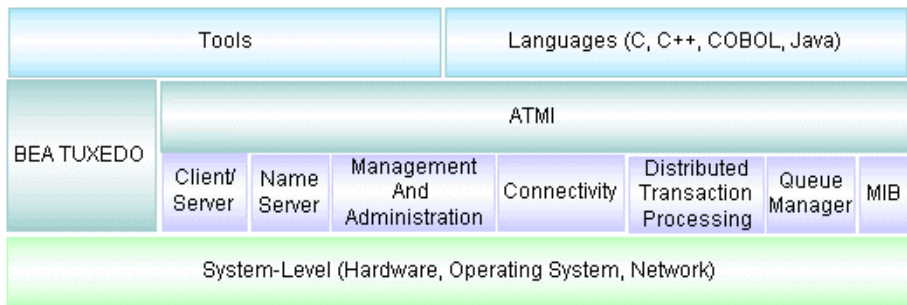
## See Also

- “BEA Tuxedo Administrative Services” on page 2-44
- “BEA Tuxedo Application Processing Services” on page 2-31

## What You Can Do Using the ATMI

The Application to Transaction Monitor Interface (ATMI), the BEA Tuxedo API, is an interface for communications, transactions, and management of data buffers that works in all environments supported by the BEA Tuxedo system. It provides the connection between application programs and the BEA Tuxedo system. The ATMI is a simple interface for a comprehensive set of capabilities. It implements the X/Open DTP model of transaction processing.

**Figure 2-2 Using the ATMI**



The ATMI supports the following tasks:

- Client initialization
- Server naming
- System messaging
- Managing transactions
- Dispatching of services
- Managing buffers

The ATMI library offers you a variety of functions for defining and controlling global transactions in a BEA Tuxedo application. Global transactions enable you to manage exclusive units of work spanning multiple programs and resource managers in your

distributed application. All work in a single transaction is treated as a logical unit, so that if any one program cannot complete its task successfully, no work is performed by programs in the transaction. Most ATMI functions support different communication styles. These functions knit together distributed programs by enabling them to send and receive data. All ATMI functions send or receive data in typed buffers. Following is a list of ATMI functions (for C and COBOL bindings), and the tasks they perform. The functions are grouped by task.

**Table 2-1 Using the ATMI Functions**

<b>For a Task Related to</b>	<b>Use This C Function</b>	<b>Or This COBOL Function</b>	<b>To</b>
Client membership	<code>tpchkauth(3c)</code>	<code>TPCHKAUTH(3cbl)</code>	Check whether authentication is required
	<code>tpinit(3c)</code>	<code>TPINITIALIZE(3cbl)</code>	Have a client join an application
	<code>tpterm(3c)</code>	<code>TPTERM(3cbl)</code>	Have a client leave an application
Buffer management	<code>tpalloc(3c)</code>	N/A	Create a message buffer
	<code>tprealloc(3c)</code>	N/A	Resize a message buffer
	<code>tpfree(3c)</code>	N/A	Free a message buffer
	<code>tpypes(3c)</code>	N/A	Get a message type and subtype
Message priority	<code>tpgprio(3c)</code>	<code>TPGPPIO(3cbl)</code>	Get the priority of the last request
	<code>tpsprio(3c)</code>	<code>TPSPRIO(3cbl)</code>	Set the priority of the next request
Request/Response communications	<code>tpcall(3c)</code>	<code>TPCALL(3cbl)</code>	Initiate a synchronous request/response to a service
	<code>tpacall(3c)</code>	<code>TPACALL(3cbl)</code>	Initiate an asynchronous request (fanout)
	<code>tpgetrply(3c)</code>	<code>TPGETRPLY(3cbl)</code>	Receive an asynchronous response
	<code>tpcancel(3c)</code>	<code>TPCANCEL(3cbl)</code>	Cancel an asynchronous request

## 2 BEA Tuxedo System Architecture

---

<b>For a Task Related to</b>	<b>Use This C Function</b>	<b>Or This COBOL Function</b>	<b>To</b>
Conversational communications	<code>tpconnect(3c)</code>	<code>TPCONNECT(3cbl)</code>	Begin a conversation with a service
	<code>tpdiscon(3c)</code>	<code>TPDISCON(3cbl)</code>	Abnormally terminate a conversation
	<code>tpsend(3c)</code>	<code>TPSEND(3cbl)</code>	Send a message in a conversation
	<code>tprecv(3c)</code>	<code>TPRECV(3cbl)</code>	Receive a message in a conversation
Reliable queuing	<code>tpenqueue(3c)</code>	<code>TPENQUEUE(3cbl)</code>	Enqueue a message to a message queue
	<code>tpdequeue(3c)</code>	<code>TPDEQUEUE(3cbl)</code>	Dequeue a message from a message queue
Event-based communications	<code>tpnotify(3c)</code>	<code>TPNOTIFY(3cbl)</code>	Send an unsolicited message to a client
	<code>tpbroadcast(3c)</code>	<code>TPBROADCAST(3cbl)</code>	Send messages to several clients
	<code>tpsetunsol(3c)</code>	<code>TPSETUNSOL(3cbl)</code>	Set unsolicited message call-back
	<code>tpchkunsol(3c)</code>	<code>TPCHKUNSOL(3cbl)</code>	Check the arrival of unsolicited messages
	N/A	<code>TPGETUNSOL(3cbl)</code>	Get an unsolicited message
	<code>tppost(3c)</code>	<code>TPPOST(3cbl)</code>	Post an event message
	<code>tpsubscribe(3c)</code>	<code>TPSUBSCRIBE(3cbl)</code>	Subscribe to event messages
	<code>tpunsubscribe(3c)</code>	<code>TPUNSUBSCRIBE(3cbl)</code>	Unsubscribe to event messages

For a Task Related to	Use This C Function	Or This COBOL Function	To
Transaction management	tpbegin(3c)	TPBEGIN(3cbl)	Begin a transaction
	tpcommit(3c)	TPCOMMIT(3cbl)	Commit the current transaction
	tpabort(3c)	TPABORT(3cbl)	Roll back the current transaction
	tpgetlev(3c)	TPGETLEV(3cbl)	Check whether in transaction mode
	tpsuspend(3c)	TPSUSPEND(3cbl)	Suspend the current transaction
	tpresume(3c)	TPRESUME(3cbl)	Resume a transaction
Service entry and return	tpsvrinit(3c)	TPSVRINIT(3cbl)	Initialize a server
	tpsvrdone(3c)	TPSVRDONE(3cbl)	Terminate a server
	tpservice(3c)	N/A	Prototype for a service entry point
	N/A	TPSVCSTART(3cbl)	Get service information
	tpreturn(3c)	TPRETURN(3cbl)	End a service function
	tpforward(3c)	TPFORWAR(3cbl)	Forward request
Dynamic advertisement	tpadvertise(3c)	TPADVERTISE(3cbl)	Advertise a service name
	tpunadvertise(3c)	TPUNADVERTISE(3cbl)	Unadvertise a service name
Resource management	tpopen(3c)	TPOPEN(3cbl)	Open a resource manager
	tpclose(3c)	TPCLOSE(3cbl)	Close a resource manager

**Note:** The use of ATMI transaction management functions is optional.

## See Also

- “Using the ATMI to Handle System and Application Errors” on page 2-28 in *Administering a BEA Tuxedo Application at Run Time*

# What Are the BEA Tuxedo Messaging Paradigms

The following table describes the BEA Tuxedo messaging paradigms available to application developers.

**Table 2-2 BEA Tuxedo Messaging Paradigms**

<b>BEA Tuxedo Messaging Paradigm</b>	<b>Description</b>
<a href="#">Conversational communication</a>	Service request mode involving multiple 2-way interactions between a client and a dedicated server
<a href="#">Event-based communication</a>	Publish/subscribe mode
<a href="#">Queue-based communication</a>	Guaranteed delivery mode
<a href="#">Request/reply communication</a>	Service request mode that can be synchronous (processing waits until the requester receives the response) or asynchronous (processing continues while the requester waits for the response)
<a href="#">Unsolicited messaging</a>	Communication from any client or server to any clients that were not requested or expected by those clients

## See Also

- “What Is Conversational Communication” on page 2-9
- “How the EventBroker Works” on page 2-10
- “What Is Queue-based Communication” on page 2-13
- “What Is Request/Reply Communication” on page 2-14
- “What Is Unsolicited Communication” on page 2-17
- “What Are Nested and Forwarded Service Requests” on page 2-18


# What Is Conversational Communication

Conversational communication is the BEA Tuxedo system implementation of a human-like paradigm for exchanging messages between clients and servers. In this form of communication, a virtual connection is maintained between the client and server. Just as in a conversation between two people, a number of messages pass back and forth between the two entities until a conclusion is reached. Over the course of the communication, both sides “remember” the point (or state) of the conversation so that relatively long operations, such as ad hoc queries, reports, and file transfers, can be supported. Conversational servers are available by default, but more can be spawned automatically if needed.

The BEA Tuxedo system provides an Application Programming Interface (API) that can be used to create conversations in applications; specifically to connect clients to servers, to send and receive messages, and to end the conversation.

Conversations can be nested but performance may be degraded as a result of doing so. Conversations may contain either transactions or service requests as appropriate. Although a conversational service can make service calls and establish conversations, those service calls and conversations cannot be forwarded. A conversation can be within the scope of, and controlled by a transaction.

Figure 2-3 Conversational Communication



## See Also


- “Using Conversational Communication” on page 1-11 in *Tutorials for Developing a BEA Tuxedo Application*

## How the EventBroker Works

The BEA Tuxedo EventBroker provides a communication paradigm in which an arbitrary number of suppliers can post messages for an arbitrary number of subscribers. Because client and server processes that use the EventBroker communicate with one another based on a set of *subscriptions*, this paradigm is known as *publish-and-subscribe* communication. The EventBroker acts like a newspaper delivery person who delivers newspapers only to customers who have paid for a subscription.



Figure 2-4 Posting and Subscribing to an Event



Event generators (either clients or servers) inform the EventBroker of changes and problems as they occur. This process is called posting an event. The EventBroker then matches the name of the event to an event name associated with a list of subscribers, and notifies each subscriber on the list of the event.

## See Also

- “What Types of Events Are Reported” on page 2-11
- “How Are Events Reported” on page 2-12
- “Using Event-based Communication” on page 1-14 in *Tutorials for Developing a BEA Tuxedo Application*

# What Types of Events Are Reported

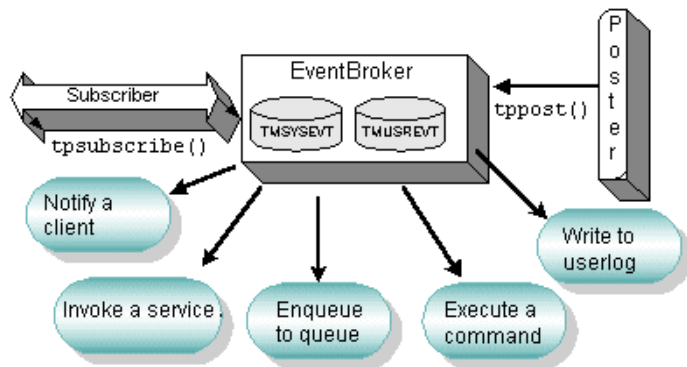
The BEA Tuxedo system supports two different types of event reports:

- System Event reports—provide details about BEA Tuxedo system events, such as servers dying, and network failures. When an event is posted by clients or servers, EventBroker matches the posted event’s name to subscriber’s of the same events and takes appropriate action determined by each subscription.
- Reports of User Events or Application Defined Events—allow application programs to post events when certain criteria are met. A banking application, for example, might post an event for withdrawals over a certain limit.

# How Are Events Reported

The EventBroker provides publish-and-subscribe functionality. A process registers a subscription with the EventBroker, indicating interest in a particular event. Subsequently, whenever the EventBroker is notified by another process that the specified event has occurred, the EventBroker reports the occurrence to any process that has subscribed for this event.

**Figure 2-5 Event-based Messaging**



The EventBroker uses several mechanisms for publishing (that is, issuing notices of) events:


- Disk-based queuing
- Asynchronous service calls
- Userlog entries
- Unsolicited messages
- System commands

# What Is Queue-based Communication

The BEA Tuxedo system offers a queue-based architecture known as /Q for applications that require persistent storage of data. The /Q component allows any client or server to store messages or service requests in queues and guarantees that any stored request is sent through the transaction protocol to ensure safe storage.

BEA Tuxedo system queues can be ordered as Last In First Out (LIFO) or First In First Out (FIFO), or on the basis of time or priority. A collection of queues is administered and referred to as a single entity known as a queue space.

**Figure 2-6 Queue-based Messaging**



## Using Application Queues

Application queues are appropriate if you must communicate in a time-independent fashion. Time-independence is a characteristic of programs that operate independently from one another and do not need to synchronize their communications simultaneously. Time-independent programs synchronize by leaving messages for each other in application queues. Messages can be dequeued in any of several ordering schemes, such as first-in, first-out (FIFO) order, priority order, or time-based order. BEA Tuxedo client and server programs can enqueue messages and dequeue messages from queues. More than one client and server can access the same queue.

To use an application queue, your program must name the queue to be accessed and the queue space in which it resides. Your application can use more than one queue space and each space can contain more than one message queue.

Because application queues reside on a disk, the availability of stored messages is guaranteed even after machine failures. To determine when the use of application queues is appropriate, you need to determine when time-independent synchronization occurs in your business, for example, in filling orders. Orders can be enqueued to disk and depending on specific order criteria, such as items or shipment location, placed in different queue spaces. Within each queue space, you can determine additional criteria, such as cost, state, and so on.

## See Also

- “Using Queue-based Communication” on page 1-15 in *Tutorials for Developing a BEA Tuxedo Application*

# What Is Request/Reply Communication


To implement request/reply communication, the BEA Tuxedo system uses IPC message queues. Queues are the key to connectionless communication. Each server is assigned an Inter-Process Communication (IPC) message queue called a request queue and each client is assigned a reply queue. Therefore, rather than establishing and maintaining a connection with a server, a client application can send requests to the server by putting those requests on the server’s queue, and then check and retrieve messages from the server by pulling messages from its own reply queue.

The request/reply model is used for both synchronous and asynchronous service requests as described in the following topics.

## What Is Synchronous Messaging

In a synchronous call, a client sends a request to a server, which performs the requested action while the client waits. The server then sends the reply to the client, which receives the reply.


**Figure 2-7 Synchronous Request/Reply Communication**



## What Is Asynchronous Messaging

In an asynchronous call, the BEA Tuxedo client does not wait for a service request it has submitted to finish before undertaking other tasks. Instead, after issuing a request, the client performs additional tasks (which may include issuing more requests). When a reply to the first request is available, the client retrieves it.

**Figure 2-8 Asynchronous Request/Reply Communication**



## See Also

- “Using the Request/Response Model (Synchronous Calls)” on page 1-7 in *Tutorials for Developing a BEA Tuxedo Application*

# What Is Unsolicited Communication


The BEA Tuxedo system offers a powerful communication paradigm called *unsolicited notification*. When unsolicited notification occurs, a BEA Tuxedo client receives a message that it has never requested. This capability makes it possible for application clients to receive notification of application-specific events as they occur, without having to request notification explicitly in real time.

Unsolicited messages can be sent to client processes by name (`tpbroadcast`) or by an identifier received with a previously processed message (`tpnotify`). Messages sent via `tpbroadcast` can originate either in a service or in another client. You can target a narrow or wide audience. You can send a message with or without guaranteed delivery to an individual client through *point-to-point notification* (`tpnotify`), or you can send information to a group of clients (`tpbroadcast`). For example, a server may alert a single client that the account about which the client is inquiring has been closed. Or, a server may send a message to all the clients on a machine to remind the users that the machine will be shut down for maintenance at a specific time.

Any process that wants to be notified about a particular event (such as a machine being shut down for maintenance) can *register* a request, with the system, to be notified automatically. Once registered, a client or server is informed whenever the specified event occurs. This type of automatic communication about an event is called *unsolicited notification*.

Because there is no limit to the number of clients and servers that may generate events and receive unsolicited notification about such events, the task of managing this category of communication can become complex. The BEA Tuxedo system offers a tool for managing unsolicited notification called the *EventBroker*.

Figure 2-9 Unsolicited Notification Messaging



## See Also

- “Using Unsolicited Notification” on page 1-13 in *Tutorials for Developing a BEA Tuxedo Application*


# What Are Nested and Forwarded Service Requests

## Nested Requests

A powerful feature of the BEA Tuxedo system is that it allows services to act as clients and call other services. Nesting is limited to two levels, which works particularly well in a 3-tier client/server architecture, that is, a system that comprises a presentation logic layer, a business logic layer, and a database layer. In such a system, the presentation layer is used to formulate a request for a particular business function that involves one or more queries to a database. Because nesting is limited to two levels, it does not degrade performance.



Figure 2-10 Nested Service Requests



## Benefit of Nested Requests

One benefit of using nested requests is that doing so enables you to keep your code small and reusable, such that each piece performs a limited task. However, if the services in your system are distributed across several servers, nested requests can lead to poor performance. While a nested request is being processed, the original service (that is, the service that issued the nested request) must wait for a response before continuing. Until a response is received, the original service cannot process another request. As a result, messages can get backed up in the request queue for the server on which this service resides.

## Example of a Nested Service Request


A customer uses a cash machine to transfer money from her savings account to her checking account. A BEA Tuxedo application performs the work necessary to transfer the money. First, on behalf of the customer, the client issues a request for a service called `TRANSFER`, and the request is placed on a queue for a server that provides that

service. Next, the TRANSFER service requests two other services, WITHDRAW and DEPOSIT, which are processed by a second server. The WITHDRAW and DEPOSIT services return responses to the TRANSFER service. Finally, TRANSFER sends a response to the client's response queue. When the client retrieves the response from the queue, the system displays a message on the screen of the cash machine, notifying the customer that the transfer is complete.

## Forwarded Requests

One alternative to nesting service requests is called request forwarding. Instead of processing a client's request, a service can pass the request to another service. The second service, also, can either process the request or pass it to another service.

**Figure 2-11 Forwarded Service Requests**



There is no limit to the number of times a request can be forwarded. Because a service that forwards a request does not need to wait for a reply from the service receiving the request, forwarding, unlike nesting requests, does not block servers. Forwarding, however, is not supported by the X/OPEN protocol X/ATMI, which may be a problem in some applications.

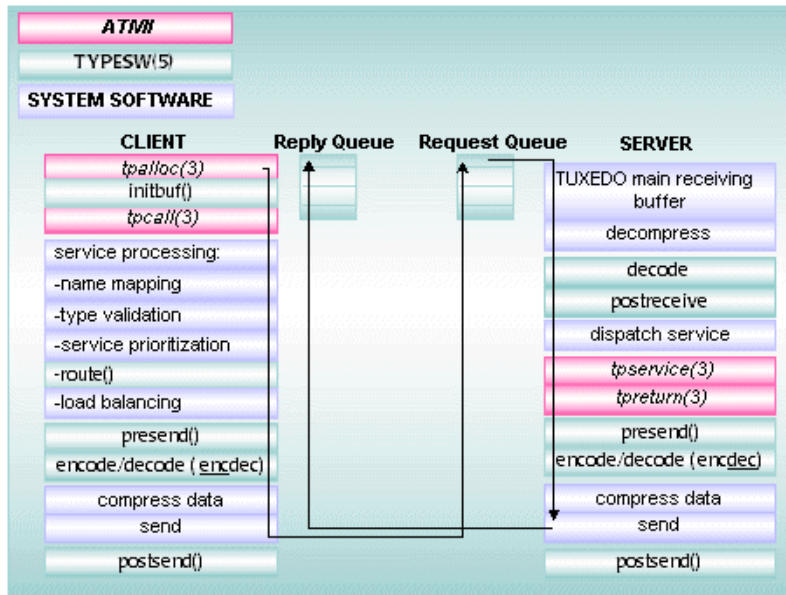
### See Also

- “Using Forwarded Calls” on page 1-10 in *Tutorials for Developing a BEA Tuxedo Application*
- “Using Nested Calls” on page 1-9 in *Tutorials for Developing a BEA Tuxedo Application*

## How the BEA Tuxedo System Processes Messages

All communication within the BEA Tuxedo system is accomplished by transferring messages. The BEA Tuxedo system passes service request messages between clients and servers through operating system Inter-Process Communications (IPC) message queues. System messages and data are passed between operating system-supported memory-based queues of clients and servers in *buffers*. In the BEA Tuxedo system, messages are packaged in *typed buffers*, buffers that contain both message data and data identifying the types of message data being sent.

Figure 2-12 Processing a Request



A client uses an ATMI function to request a service by name. A *naming* facility is used to check the MIB to determine whether the specified service is currently available. The BEA Tuxedo system uses an automatic routing option to map messages that meet specific criteria (message value) to a specific server. This is called *data-dependent routing*. If messages use data-dependent routing, the system uses the data in the buffer for the routing algorithm. This algorithm provides a method of selecting a group of servers that can process the service request. To avoid burdening a few servers with many requests while leaving other servers that advertise the same services idle, the BEA Tuxedo system maintains a set of metrics in the MIB that help it distribute service requests evenly across all servers. This practice is called *load balancing*.

A local service request may be prepared for a selected server and enqueued on that server's queue with a predefined priority. This practice is called *service prioritization*. Once the service request is on the server, the run-time system retrieves the message in priority order. The message is dispatched to the appropriate service and processed. Then the results are returned to the client queue.

BEA Tuxedo system-provided software offers features that an application can automatically and routinely use during message processing. These features include: data encoding and decoding, data compression and decompression, transactional

context setting, and security processing, to name a few. In addition, the BEA Tuxedo system software invokes application business logic by dispatching a service function and passing it to the appropriately preprocessed buffer.

The service routine is executed and returns a reply (also a typed buffer). The run-time system prepares the reply for the client by encoding the message automatically: it packages the data in such a way that it can be transmitted between machines on which different types of byte ordering are used, allowing data to cross network and platform boundaries. The system then sends the message to the client. This process is called *data encoding*. The run-time system on the client retrieves the reply message, decodes it if necessary, and delivers the FML buffers (or buffers of another message buffer type) to package the application data. Type validation, encoding, routing, and load balancing are performed as required. Service requests can be performed synchronously or asynchronously.

Remote requests travel through the local BRIDGE to the remote machine, where the remote BRIDGE simply acts as a client and the request is processed as if the client and server were on the same machine. The BRIDGE provides standard data encoding/decoding and uses standard network transports to communicate. BRIDGEs look like ordinary local servers to clients and servers.

## What Are the Benefits of Service Request Processing

- Connectionless processing—This processing, coupled with direct client/server communication, reduces the overhead associated with establishing a connection.
- Reduced network traffic—Service requests invoke potentially complex services on remote machines, sending only the minimum data required and receiving minimal results.

## See Also

- “What Are the BEA Tuxedo Messaging Paradigms” on page 2-8
- “What Are Typed Buffers” on page 2-24

## What Are Typed Buffers

All ATMI functions send or receive data using typed buffers. The BEA Tuxedo system handles translations and data conversions between dissimilar machines. By using buffers, BEA Tuxedo programs avoid the need to translate data that crosses different platforms with different data representations.

A buffer is a memory area that serves as a logical container for data. When a buffer contains no metadata (that is, no information about itself), then it is an *untyped buffer*. When a buffer includes metadata such as information that can be stored in it (for example, a type and subtype, or string names that characterize a buffer), then it is a *typed buffer*.

Typed buffers can be transmitted over any network, on any operating system, with any protocol supported by the BEA Tuxedo system. They can also be used on platforms with different data representations. As a result, the use of typed buffers facilitates the tasks of translation and data conversion between dissimilar machines.

The BEA Tuxedo system supports five sorts of typed buffers:

- `STRING`
- `VIEW`
- `CARRAY`
- `FML`
- `XML`

You assign buffer types in the `ENVFILE` parameter defined in the `MACHINES` section of the configuration file. Assigning or overriding them in the `ENVFILE` parameter in the `SERVERS` section of the configuration file can make them unavailable to processes that require them.

Definitions of the various types of message buffers are provided in the description of `tm_typesw` in `tuxtypes(5)` in *BEA Tuxedo File Formats and Data Descriptions Reference*. It is to your advantage to change `tm_typesw` so it contains only buffer types specifically needed by a given server.

## Characteristics of Buffer Types

When you use ATMI communication functions, your application must first use `tpalloc` to get a buffer from the system, specifying its size, type, and optionally subtype. The BEA Tuxedo system recognizes and processes the buffer type, so that your data is transmitted over any type of network, protocol, and operating system supported by the BEA Tuxedo system. The following table describes the different types of buffers available in a BEA Tuxedo environment.

**Table 2-3 Buffer Types Characteristics**

<b>This Typed Buffer</b>	<b>Is Defined As Follows</b>	<b>And Is Used for</b>
CARRAY	<p>Character array type is a collection of characters that is handled opaquely:</p> <ul style="list-style-type: none"><li>■ Characters are not interpreted in any way.</li><li>■ No subtypes are specified.</li><li>■ Your application must specify the buffer length for CARRAY message buffers used as input to ATMI functions.</li></ul>	<p>Data that will not be interpreted by the BEA Tuxedo system and for which data-dependent routing, encoding, or decoding is not required.</p>



This Typed Buffer	Is Defined As Follows	And Is Used for
FML	<p>Field Manipulation Language (FML) is a data structure that stores tagged values. Values are typed, may be specified more than once, and vary in length.</p> <p>The FML buffer is an abstract data type used in operations to create, modify, delete, or access fields. In your program, you access or update a field in the fielded buffer by referencing the identifier, and the FML function provides for a run-time translation of the field's location and data type, and performs the operation.</p> <p>One interface to FML uses 16 bits (FML16) for field identifiers and lengths of fields; the other uses 32 bits (FML32).</p> <ul style="list-style-type: none"> <li>■ The 16-bit version allows for up to approximately 8000 unique fields, character strings, and arrays of up to 64,000 bytes, and similar lengths for the entire buffer.</li> <li>■ The 32-bit interface allows for millions of unique fields and buffer lengths of up to two billion bytes.</li> </ul> <p>The functionality of the two interfaces is identical. The power of FML is in its flexibility. The size of the buffer can vary, depending on the needs of the application for each message. Character fields may also vary in length, so wasted space is avoided.</p> <p>Fielded buffers offer data independence to the application. When writing an application, you do not need to know how or where the data is stored within a fielded buffer. FML provides associative field access, so you simply specify a field by name and its value is returned. FML also contains conversion functions, so that you can store or retrieve a field in a particular data format, regardless of the underlying storage type.</p> <p>FML buffers also support storage of more than one value for a field. The variable length format of fielded buffers allows for multiple field occurrences to be stored and retrieved.</p> <p>Fielded buffers provide a convenient way to transfer a collection of fields, perhaps different with each message, from a client to a server and back, or to store fields in an application queue. We recommend using FML, particularly if the interface between clients and servers may change.</p>	<ul style="list-style-type: none"> <li>■ Communications</li> <li>■ Creating, modifying, deleting, or accessing fields during operations</li> </ul>
STRING	<p>A set of non-null characters ending with a null character. The data type is character and the length is determined by counting characters in the buffer until reaching the null character. No subtype is specified.</p>	C programs

<b>This Typed Buffer</b>	<b>Is Defined As Follows</b>	<b>And Is Used for</b>
VIEW	<p>A VIEW is simply a C structure or a COBOL record that has an associated definition of which fields and their types appear in the record in which order. This buffer is used for fixed collections of data elements, or structures or records; its subtype is used to specify the record format name.</p> <p>VIEW records are flat data structures. They do not support structures within other structures, nor do they allow arrays of structures or pointers. They support integral data types such as long integer, character, and decimal.</p> <p>VIEWS are provided as a way to use C structures and COBOL records with the BEA Tuxedo system. The BEA Tuxedo run-time system understands the record format based on the view description read at run time. When allocating a VIEW, your application specifies a buffer type of VIEW and a subtype that matches the name of the view. The run-time system can do the following:</p> <ul style="list-style-type: none"><li>■ Determine how much space is needed, based on structure size, so the application need not specify buffer length</li><li>■ Compute how much data to send in a request or response, and handle encoding and decoding when a message is transferred between different machine types</li></ul>	C structures and COBOL records used with a BEA Tuxedo application
XML (Extensible Markup Language)	<p>XML buffers enable BEA Tuxedo applications to use XML for exchanging data within and between applications. BEA Tuxedo applications can send and receive simple XML buffers, and route those buffers to the appropriate servers. All logic for dealing with the XML documents, including parsing, resides in the application. An XML document consists of: a sequence of characters that encode the text of a document and a logical structure of the document and meta-information related to the structure.</p> <p>The XML parser in the BEA Tuxedo system performs autodetection of character encodings, character code conversion, detection of element content and attribute values, and data type conversion.</p> <p>Data-dependent routing is supported for XML buffers.</p>	<ul style="list-style-type: none"><li>■ XML documents and datagrams</li><li>■ Data interchange between humans and machines, such as from a Web server to a user's browser</li><li>■ Data exchange between applications, or from machine to machine</li></ul>

## See Also

- “Customizing a Buffer” on page 3-28 in *Programming a BEA Tuxedo Application Using C*

## Using the MIB

The MIB programming interface enables you to manage operations in the BEA Tuxedo system easily. Specifically, it allows you to monitor, configure, and tune your application through your own programs. The MIB can be defined as:

- An implementation-independent management database defined as a set of FML attributes
- A programming interface that enables you to query the BEA Tuxedo system (that is, to obtain information from the system through a `get` operation) or to update the BEA Tuxedo system (that is, to change information in the system through a `set` operation) at any time using a set of ATMI functions. Examples of these functions include `tpalloc`, `tprealloc`, `tpgetrply`, `tpcall`, `tpacall`, `tpenqueue`, and `tpdequeue`.

## See Also

- MIB(5) in *BEA Tuxedo File Formats and Data Descriptions Reference*
- “Types of MIB Users” on page 2-30
- “Classes, Attributes, and States in the MIB” on page 2-30

## Types of MIB Users

The MIB defines three types of users: system administrators, system operators, and others. The following table describes each type.

Type of User	Characteristics
Application administrator	Person responsible for keeping an application running successfully. The administrator is authorized to use all administrative tools and all MIB administrative capabilities. The administrator configures, manages, and modifies a running production application.
System operator	Monitors and reacts to the daily operation of a production application. The operator monitors statistics about a running application, sometimes reacting to events and alerts by taking actions such as booting servers or shutting down machines. An operator does not reconfigure an application, add servers or machines, or delete machines.
Other	People or processes (such as custom programs) that may need to read the MIB but are not authorized to change the application.

## Classes, Attributes, and States in the MIB

*Classes* are the types of entities, such as servers and machines, that make up a BEA Tuxedo application. *Attributes* are characteristics of the objects in a class: identity, state, configuration parameters, run-time statistics, and so on. There are a number of attributes that are common to MIB operations and replies, and common to individual classes. Every class has a *state* attribute that indicates the state of the object. The state of an object is either *return to the user* or *new, changed state*, if you are invoking an operation on the MIB to change an object's state.

Independent of classes is a set of common attributes that are defined in the MIB(5) reference page. These attributes control the input operations, communicate to the MIB what the user is trying to do, and/or identify to the programmer some of the characteristics of the output buffer that are independent of a particular class.

## BEA Tuxedo Application Processing Services

The BEA Tuxedo system offers the following application processing services:

- [Data compression](#)
- [Data-dependent routing](#)
- [Data encoding](#)
- [Data encryption](#)
- [Data marshalling](#)
- [Load balancing](#)
- [Message prioritization](#)
- [Service and event naming](#)

## What Is Data Compression


*Data compression* is the process of shrinking an application buffer so it can be transmitted more quickly across a network or to a remote domain. By setting a maximum size for an application buffer, you can make sure that compression is triggered automatically for application buffers that match or exceed a specified size. When the buffer arrives at its destination, its data is decompressed, that is, restored to its original size.

Data compression, performed before files are shipped between machines, improves network performance.

The process of compression enhances security slightly because it involves scrambling the data.

**Note:** Data compression also occurs frequently during encryption.

**Figure 2-13 Data Compression**



## What Is Data-dependent Routing

The BEA Tuxedo system uses an operation called data-dependent routing to enable a client to send requests for the same service to multiple copies of that service. Which copy of the service eventually accepts and processes the request is determined by the data in the request message. Once an administrator has set up data-dependent routing for an application, client requests can be routed automatically to servers based on the data in the requests.

When an application includes multiple copies of the same service, each copy is assigned a unique purpose, just as the first volume of a multivolume encyclopedia contains entries that begin with the letter “A.” A list of all copies of the service, along with identifying information about the purpose of each, is kept in a set of routing tables in the BEA Tuxedo bulletin board. When the system receives a client request, it finds an identifying string in the request message and searches the routing tables in the bulletin board for the same string. On the basis of this match, the system identifies the appropriate server to which it can forward the client request.

**Note:** The bulletin board routing tables can be modified as necessary.

## Uses of Data-dependent Routing

Data-dependent routing is useful when clients issue service requests to:

- Horizontally partitioned databases
- Rule-based servers
- Distributed Application

A horizontally partitioned database is an information repository that has been divided into segments, each of which is used to store a different category of information. This arrangement is similar to a library in which each shelf of a bookcase holds books for a different category (for example, biography, fiction, and so on).

A rule-based server is a server that determines whether service requests meet certain, application-specific criteria before forwarding them to service routines. Rule-based servers are useful when you want to handle requests that are almost identical by taking slightly different actions for business reasons.

A distributed application consists of one or more local or remote clients that communicate with one or more servers on several machines linked through a network. A client (or server acting as a client) issues a request for a particular service. The *address* of the request is determined by data (carried in the same buffer that conveys the request), identifying the server that can fulfill the request. More than one server may be able to do so. The BEA Tuxedo system selects a server to receive the request by matching the data to the routing criteria provided in the bulletin board.


## Example of Data-dependent Routing with a Horizontally Partitioned Database

Suppose two clients in a banking application issue requests for the current balance in two accounts: Account 3 and Account 17. If data-dependent routing is being used in the application, then the BEA Tuxedo system performs the following actions:

1. Gets the account numbers for the two service requests (3 and 17).
2. Checks the routing tables on the BEA TUXEXDO bulletin board that show which servers handle which range of data. (In this example, server 1 handles all requests for accounts 1 through 10; server 2 handles all requests for accounts 11 through 20.)
3. Sends each request to the appropriate server. Specifically, the system forwards the request about Account 3 to server 1, and the request about account 17 to server 2.

The following figure illustrates this process.

**Figure 2-14 Data-dependent Routing with a Horizontally Partitioned Database**





## Example of Data-dependent Routing with Rule-based Servers

A banking application includes the following rules:


- Customers can withdraw up to \$500 without entering a special password.
- Customers must enter a special password to withdraw more than \$500.

Two clients issue withdrawal requests: one for \$100 and one for \$800. If data-dependent routing is enabled to support the withdrawal rules, then the BEA Tuxedo system performs the following actions:

1. Gets the amount specified for withdrawal in the two service requests (\$100 and \$800).
2. Checks the routing tables on the BEA Tuxedo bulletin board that show which servers handle request for the amount being requested. (In this example, server 1 handles all requests to withdraw amounts up to \$500; server 2 handles all requests to withdraw amount over \$500.)
3. Sends each request to the appropriate server. Specifically, the system forwards the request for \$100 to server 1 and the request for \$800 to server 2.

The following figure illustrates this process.

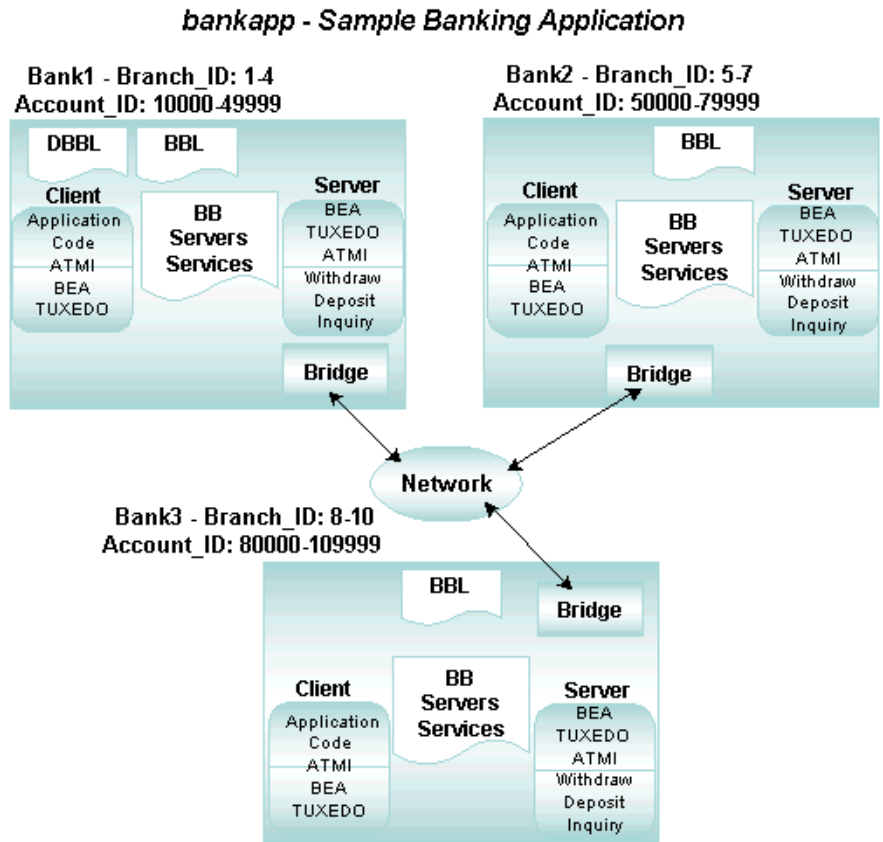
**Figure 2-15 Data-dependent Routing with Rule-Based Servers**



## Example of Data-dependent Routing with Distributed Application

The following diagram shows how client requests are routed to servers. In this example, a banking application called `bankapp` uses data-dependent routing. `bankapp` has three server groups (BANK1, BANK2, and BANK3) and two routing criteria (Account ID and Branch ID). The services WITHDRAW, DEPOSIT, and INQUIRY are routed using the Account\_ID field; the services OPEN and CLOSE are routed using the Branch\_ID field.

Figure 2-16 Sample Banking Application Using Routing Criteria



In the preceding diagram, requests are routed as indicated in the following table.

<b>Withdrawals, Deposits, Inquiries, and Openings or Closings of the Following Accounts . . .</b>	<b>Are Routed to . . .</b>
Numbers 10000—49999 for branches 1—4	Bank1
Numbers 50000—79999 for branches 5—7	Bank2
Numbers 80000—109999 for branches 8—10	Bank3

## What Are Encoding and Decoding of Data


*Encoding* and *decoding* enable messages with different data representations (for example, byte ordering or character sets) to be transferred between machines. The BEA Tuxedo system accomplishes this by encoding and decoding data to a machine-independent representation for transmission. It employs, by default, the XDR algorithm, which can be customized by replacing the BEA Tuxedo system functions with user-written functions. Encoding and decoding are used only between machines and only when a remote machine uses a data representation other than the one used on the local machine. Encoding and decoding allow machines with different data architectures to operate within a heterogeneous BEA Tuxedo system. Programmers can manage data in representations natural to their own environments.

The BEA Tuxedo system uses buffer types to determine the type of fields contained in a message, and to perform the mapping required for coding tasks. This mapping is not performed by unstructured buffer types such as `X_OCTET` and `CARRAY`. Thus, developers using `X_OCTET` and `CARRAY` buffers are free to deploy in mixed-machine environments.

## What Is Data Encryption

*Encryption* is the act of converting a message into a coded format that is unintelligible to users. When an encrypted message arrives at its destination, it is decrypted, that is, converted back to its original format.

**Figure 2-17 Data Encryption**



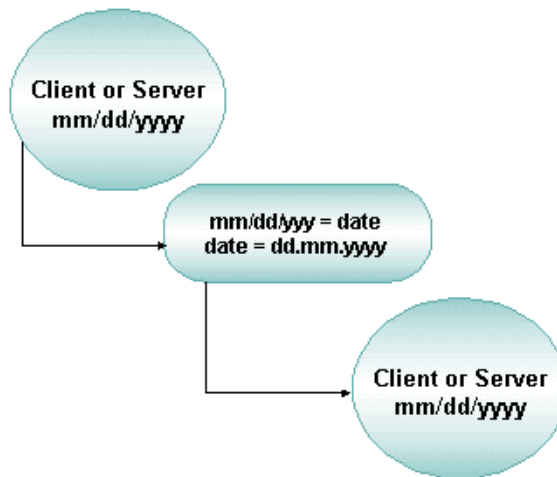
Encryption does not increase the number of bits in the data, but it adds processing time to the task of sending a message. Because data is compressed during encryption, however, lost processing time may be bought back, since less data is being sent across the network. When data is compressed, there is also a moderate boost to security, because the data is somewhat scrambled during compression.

## What Is Data Marshalling

Data marshalling is a method of handling information through the language-based TxRPC (X/Open-TxRPC) offered by the BEA Tuxedo system. TxRPC is a set of protocols for remote procedure calls that supports global transactions. Though a TxRPC call looks like a local procedure call, when a C function is called, the arguments passed to the function are packaged so they can be sent to a server that performs the work of the called function. This argument packaging is called *marshalling*. A function's arguments are *marshalled* or packaged in a way that allows them to cross network and platform boundaries, and then *unmarshalled* at their destination before being passed to the invoked remote procedure, ready for use.

This process is transparent to the client (the calling program) and the server (the remote procedure). The marshalling and unmarshalling routines are generated automatically by the BEA Tuxedo Interface Definition Language (IDL) compiler. An IDL compiler takes a description of a set of RPCs and generates routines, called stubs, for the client and server programs. These stubs contain marshalling and unmarshalling logic, as well as the communication logic that allows a client and server to exchange marshalled data.

**Figure 2-18 Data Marshalling**



## What Is Load Balancing

*Load balancing* is a technique used by the BEA Tuxedo system for distributing service requests evenly among servers that offer the same service. This avoids overburdening some servers while leaving others idle or infrequently used. Before sending a request to a service routine, the system identifies all servers capable of handling the request and selects the one most appropriate for maintaining a balanced load across all the servers in the configuration.

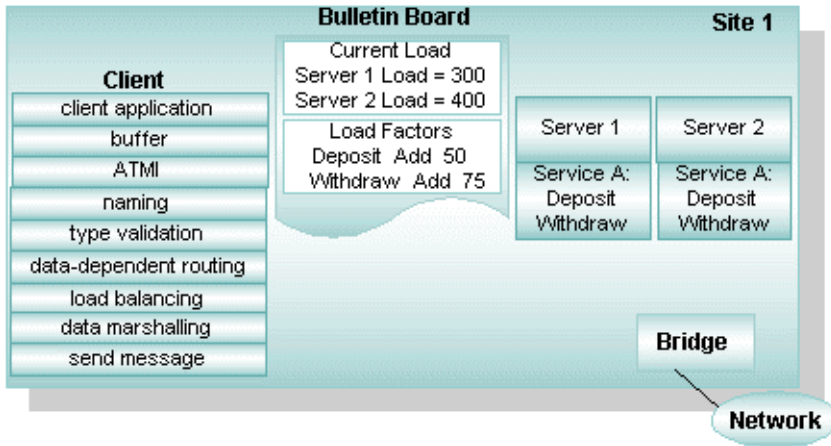
# Assigning a Load Factor

*Load* refers to a number assigned to a service request based on the amount of time required to execute that service. Loads are assigned to services so that the BEA Tuxedo system can understand the relationship between requests. To keep track of the amount of work, or total load, being performed by each server in a configuration, the administrator assigns a *load factor* to every service and service request. A load factor is a number indicating the amount of time needed to execute a service or a request. On the basis of these numbers, statistics are generated for each server and maintained on the bulletin board on each machine. Each bulletin board keeps track of the cumulative load associated with each server, so that when all servers are busy, the BEA Tuxedo system can select the one with the lightest load.

You can control whether a load-balancing algorithm is used on the system as a whole. Such an algorithm should be used only when necessary, that is, only when a service is offered by servers that use more than one queue. Services offered by only one server, or by multiple servers in an MSSQ (multiple server single queue) do not need load balancing. The `LDBAL` parameter for these services should be set to `N`. In other cases, you may want to set `LDBAL` to `Y`.

To determine how to assign load factors (in the `SERVICES` section of `UBBCONFIG`), run an application for a long period of time and note the average time it takes to perform each service. Assign a `LOAD` value of 50 (`LOAD=50`) to any service that takes roughly the average amount of time. Any service taking longer than average should have a `LOAD>50`; any service taking less than the average should have a `LOAD<50`.

Figure 2-19 Load Balancing

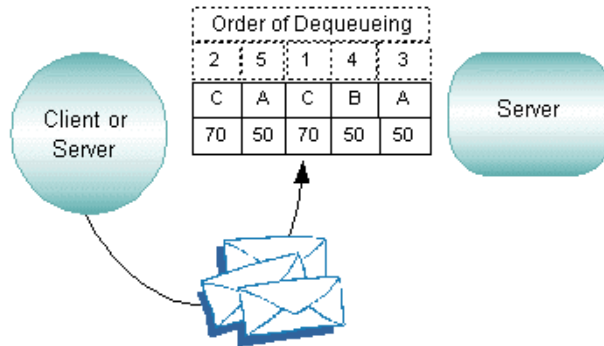


## What Is Message Prioritization

Priorities determine the order in which service requests are dequeued by a server. Priority is assigned by a client to individual services and can range from 1 to 100, where 100 represents the highest priority.

All services are assigned a starting priority of 50. A server's starting priority can be changed during application configuration. Once you have defined your set of services, you can assign the appropriate priorities to them. For example, your business may require that some services have a relatively high priority of 70, which means those services are dequeued before those with the lower priority of 50. In the following illustration, a server offers services A (with a priority of 50, B (with a priority of 50), and C (with a priority of 70).

Figure 2-20 Prioritization of Messages



A request for service C is always dequeued before a request for A or B due to the higher priority of C. Requests for A and B have equal priority. This feature is useful in applications in which not all requests are equally urgent or important.

A “starvation prevention” mechanism prevents low-priority messages from waiting endlessly on the queue. Every tenth message is dequeued in FIFO (first in first out) order regardless of priority; the first through the ninth messages are dequeued in order of priority.

## What Is Meant by Naming

The BEA Tuxedo system uses three naming devices: service names, message queue names, and event names. Names can be any words or alphanumeric strings, as long as they do not begin with a period (“.”). Because administrative servers use the BEA Tuxedo system infrastructure, system and application resources must be clearly distinguished.

## Naming Services


When services are named, an application component can locate another component through a name. Names can be simple words (such as “deposit”) or alphanumeric strings (such as “deposit2”). Names should be selected on the basis of the scope of the



application and a map that contains the global picture of the relationships among application components. These maps or services are like the pages in a telephone book for application components.

When a BEA Tuxedo system server is activated, the bulletin board (the dynamic part of the MIB) advertises the names of its services. Service names are associated with a server's physical address so that requests can be routed to that server. Names that programmers use in their applications are completely location transparent. When a client program asks for a service by name, the BEA Tuxedo system consults its name registry in the bulletin board. The name registry provides the information necessary to convert the string name (for example, `TICKET`) to a machine name and the physical address of a server that advertises that service. The BEA Tuxedo system then sends the request to the appropriate server.

Figure 2-21 Locating a Service by Name



## Advertising Services

The BEA Tuxedo system uses two administrative servers to coordinate the distribution of information on the bulletin board to all active machines in the application:

- **DBBL**—The Distinguished Bulletin Board Liaison server propagates global changes to the MIB and maintains the static part of the MIB. The DBBL coordinates the state of different machines involved in an application. Only one DBBL exists for an entire application. It can be migrated to other machines for fault resiliency.
- **BBL**—The Bulletin Board Liaison server maintains the bulletin board. A BBL resides on every active machine in an application. The BBL coordinates changes to the local MIB and verifies the integrity of application programs active on its machine.

# Naming Events

The BEA Tuxedo system offers a publish-and-subscribe mechanism: clients and servers can dynamically register or unregister a standing request to receive alerts (or messages) when a particular event occurs. Other clients and servers post user-defined or system events as they occur in the application. When a client or server no longer needs to be notified about a particular event, the relevant subscription can be cancelled.

## See Also

- “How the EventBroker Works” on page 2-10

# BEA Tuxedo Administrative Services

A set of system servers provides the following administrative services needed by the BEA Tuxedo system:

- [Application queue management](#)
- [Centralized application configuration](#)
- [Distributed application management](#)
- [Dynamic application reconfiguration](#)
- [Event management](#)
- [Security management](#)
- [Startup and shutdown of an application](#)
- [Transaction management](#)
- [Workstation management](#)

**Note:** For information on administrative services, see the topic, “Three Ways of Viewing the BEA Tuxedo System Infrastructure” on page 3-1

# 3 Three Ways of Viewing the BEA Tuxedo System Infrastructure

- Basic BEA Tuxedo System Infrastructure
- Management View: Using Administrative Tools
- BEA Tuxedo Administrative Services
- Development View: What You Can Do Using the ATMI
- Run-Time System View: Using Tools in Different Configurations

## Basic BEA Tuxedo System Infrastructure

The BEA Tuxedo system provides an infrastructure for the efficient routing, dispatching, and management of application service requests, event postings and notification, and application queues. This infrastructure can be explored from three perspectives:

- *Administrative or management perspective*—Encompasses a variety of tools available to manage your application.
- *Development (using the ATMI) perspective*—Encompasses those tasks you can perform using the ATMI. Clients request services through the ATMI. Server

programs group several services, which are invoked according to the rules defined by the ATMI. Application designers construct client and server programs by linking the BEA Tuxedo run-time system with their application code.

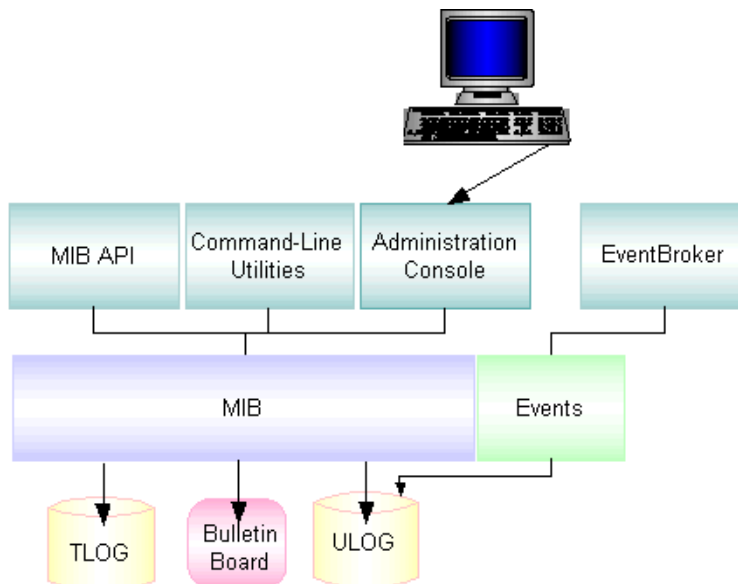
- *BEA Tuxedo run-time system view*—Encompasses single, distributed, and multiple domain configurations.

## Management View: Using Administrative Tools

The BEA Tuxedo System MIB contains all the information necessary for the operation of an application. The MIB is designed to be programmable, so that you can write custom administrative programs. Administrative tools are constructed around the MIB and provide different types of interfaces to it. These tools include the following:

- *BEA Administration Console*—a Web-based tool used to monitor an application, and to dynamically configure it.
- *BEA Tuxedo administrative servers*—servers that automate most of the management tasks for a distributed application, such as naming services and events, starting up and shutting down an application, dynamically reconfiguring an application, and so on.
- *BEA Tuxedo MIB Application Programming Interface*—a set of functions for accessing and modifying information in the MIB.
- *Command-line utilities*—a set of commands used to activate, deactivate, configure, and manage an application (that is, `tmboot`, `tmshutdown`, `tmconfig`, and `tmadmin`, respectively). (See *BEA Tuxedo Command Reference*.)
- *EventBroker*—a mechanism that informs administrators of faults or exceptional happenings.

Figure 3-1 Tools to Administer Your Application



## Available BEA Tuxedo System MIBs

The Management Information Base comprises a *core MIB*, which is common to all applications, and several component MIBs, which are optional. The core MIB, called `TM_MIB`, defines the parts of an application that are required in every BEA Tuxedo application. It is also used to administer those parts of an application. `TM_MIB` defines a BEA Tuxedo system application as a set of classes (for example, servers, groups, machines, domains), each of which is made up of objects that are characterized by various attributes (for example, identity and state).

Each of the component MIBs describes a subsystem of the BEA Tuxedo system. The following components are currently available:

- `ACL_MIB`—used to administer Access Control Lists
- `APPQ_MIB`—used to administer application stable-storage queues
- `EVENT_MIB`—used to control event notification and the subscription request database
- `WS_MIB`—used to manage Workstation groups and processes associated with them

## Using the BEA Administration Console

The BEA Administration Console is a graphical user interface to the BEA Tuxedo system, specifically to the MIB. It is accessed through the World Wide Web. The BEA Administration Console makes common administrative functions available through a browser. It allows any user with the currently available browser to administer a BEA Tuxedo application.

To use the BEA Administration Console, perform the following tasks:

1. Point your browser to the URL for a machine in your domain on which the Console server components reside.
2. Initiate a download of Java applets. The applets implement the Console and establish communication with the server.

## Browser Requirements

Each release of the BEA Tuxedo system supports the currently available browsers. Consult the following BEA Web site for information about browsers currently supported by the BEA Administration Console.

## See Also

- “Benefits of Using the BEA Administration Console” on page 3-5
- “Exploring the Main Menu of the BEA Administration Console” on page 3-6
- “Ways to Monitor Your Application” on page 2-2 in *Administering a BEA Tuxedo Application at Run Time*

# Benefits of Using the BEA Administration Console

- *Authentication*—The BEA Administration Console forces users to identify themselves. It prompts the administrator for a user name and password. This information is communicated in an encrypted fashion between the browser and the server, where the user's identity is then verified. (Much of the server setup is done during installation, when server components of the Console are installed and made available to the Web server.)
- *Context-sensitive help*—Context-sensitive help is available for all Console screens and tools. You can request information about any field or area of a screen simply by dragging a question mark icon to that field or any area and clicking.
- *Encryption*—The data transferred between the server side and the browser is compressed (56-bit or 128-bit encryption) so that no one can read it. This makes the system resistant to anyone trying to inject false administrative protocol messages into the stream.
- *Firewall readiness*—The port on which the BEA Administration Console server listens and interacts with the browser is well defined and configurable; you can configure it to match ports that you want to allow through your firewall. This capability enables you to do Console-based administration through your firewall, if necessary.
- *Icons*—The icons used in the BEA Administration Console connote state (for example, *not active*) or represent particular objects in the application, for example, machines or servers.
- *Java-capable browser*—The Java browser supports the Java virtual machine that runs the applets and enables communication.
- *No client-side installation*—No installation is required on your machine. Point your browser to the URL for a machine in your domain on which the Console server components reside. Then initiate a download of Java applets. The applets implement the Console and establish communication with the server.
- *Universal secure access*—From any Java-capable browser, you can access the system from anywhere in the world with confidence that security mechanisms are already in place.

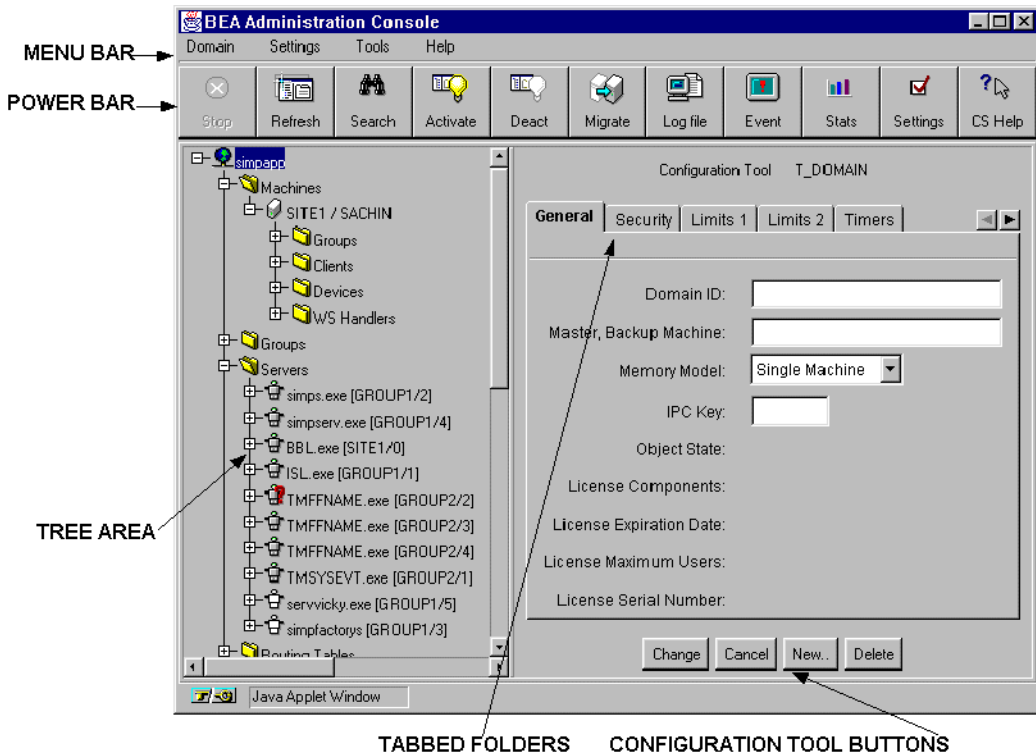
## Exploring the Main Menu of the BEA Administration Console

When you first bring up the Web and invoke the BEA Administration Console, the main window is displayed. The main window is divided into four major sections:

- *MENU BAR*—A row of frequently used menus.
- *POWER BAR*—A row of buttons that provide shortcuts to frequently used tools, such as Help.
- *TREE*—A hierarchical representation of the administrative class objects (such as servers and clients) in a BEA Tuxedo system domain.
- *CONFIGURATION TOOL*—A set of tabbed folders on which you can display, define, and modify the attributes of objects, such as the name of a machine.



Figure 3-2 Main Menu of the Administration Console



**Note:** The Power Bar buttons and some menu items are not fully displayed unless you are connected to a domain.

## Using the Configuration Tool

As soon as you start using the Configuration Tool, the right-hand column dedicated to that tool is populated with tabbed folders in which you enter information needed for configuration.

The tabbed folders in the Configuration Tool area are electronic forms that display and solicit information about the attributes of an administrative object. A set of folders is provided for each administrative class of objects (such as machines and servers). The

number of attributes associated with a class varies greatly, depending on the class. Therefore, you may find anywhere from one to eight folders displayed when you invoke the tabbed folders for a selected object.

When the Configuration Tool area is populated, another row of buttons is displayed in the main window, below the tabbed folders. These four buttons allow you to control the configuration work done in the folders.

## What Is the Tree

The Tree, which appears in the left-hand column of the main BEA Administration Console window, is a hierarchical representation of the administrative objects that make up a single BEA Tuxedo system domain. The BEA Administration Console graphically depicts the relationship between each object and the others by showing its nesting level and parent objects. You can choose to view a complete Tree (comprising all configurable objects of all types in the domain) or a subset of objects.

After you have set up and activated a domain, the Tree is populated with labeled icons, representing the administrative class objects in your domain.

## What Is an Administrative Object

The Tree consists of a list of classes. Under each class name, the Tree provides a list of the objects included in that class. For example, suppose your domain includes two machines (both at SITE1) named *romeo* and *juliet*. Both machines are listed in the Tree below the name of the class to which they belong: Machines. Thus they will be listed as follows:

```
Machines
    SITE1/romeo
    SITE1/juliet
```

The name of each object in the Tree is preceded by an icon. Each machine, for example, is represented by a computer; each client, by a human figure.

## Using the Power Bar

The Power Bar appears near the top of the main BEA Administration Console window, immediately below the Menu Bar. The Power Bar is a row of 12 buttons that allow you to invoke tools for frequently performed administrative operations. They are labeled with both icons and names. The following table explains the purpose of each button.

<b>Press</b>	<b>To</b>
Stop	Interrupt the current operation and return control to the administrator (who can then request a new operation).
Refresh	Redraw the current window to clarify text and graphics.
Search	Look for a particular object class or object in the Tree.
Activate	Begin the entire BEA Tuxedo system configuration or a selected part of it.
Deactivate	Terminate the entire BEA Tuxedo system configuration or a selected part of it.
Migrate	Move a server group or machine to an alternative location, or swap the MASTER and BACKUP machines.
Logfile	Display the ULOG file from a particular machine in the active domain.
Event Tool	Monitor system events.
Stats	View a graphical representation of BEA Tuxedo system activity.
Options	Assign parameters for your administrative BEA Administration Console session: (a) specify the location of online documentation; (b) have objects sorted by name or class; and (c) specify view-only or edit mode as your default working mode.
CS Help	Get context-sensitive help. Click on a designated area in a BEA Administration Console window. A help window opens, displaying text about the selected topic.

Press	To
Help	Go to the beginning of the Administration topics, so you can choose information that will be most helpful to you.

---

## Managing Operations Using the MIB

The AdminAPI is an application programming interface (API) for directly accessing and manipulating system settings in the BEA Tuxedo Management Information Bases (MIBs). You can use the AdminAPI to automate administrative tasks, such as monitoring log files and dynamically reconfiguring an application, thus eliminating the need for human intervention. This advantage can be crucially important in mission-critical, real-time applications. Using the MIB programming interface, you can manage operations in the BEA Tuxedo system easily. Specifically, you can monitor, configure, and tune your application through your own programs. The MIB can be defined as:

- An implementation-independent management database defined as a set of FML attributes
- A programming interface that enables you to query the BEA Tuxedo system (that is, to obtain information from the system through a `get` operation) or to update the BEA Tuxedo system (that is, to change information in the system through a `set` operation) at any time using a set of ATMI functions. Examples of these functions include `tpalloc`, `tprealloc`, `tpgetrply`, `tpcall`, `tpacall`, `tpenqueue`, and `tpdequeue`.

### See Also

- MIB(5) in *BEA Tuxedo File Formats and Data Descriptions Reference*
- “Types of MIB Users” on page 3-11
- “Classes, Attributes, and States in the MIB” on page 3-11

# Types of MIB Users

The MIB defines three types of users: system administrators, system operators, and others. The following table describes each type.

Type of User	Characteristics
Application administrator	Person responsible for keeping an application running successfully. The administrator is authorized to use all administrative tools and all MIB administrative capabilities. The administrator configures, manages, and modifies a running production application.
System operator	Person responsible for monitoring and reacting to the daily operation of a production application. An operator monitors statistics about a running application, sometimes reacting to events and alerts by taking actions such as booting servers or shutting down machines. An operator does not reconfigure an application, add servers or machines, or delete machines.
Other	People or processes (such as custom programs) that may need to read the MIB but are not authorized to change the application.

## Classes, Attributes, and States in the MIB

*Classes* are the types of entities such as servers and machines that make up a BEA Tuxedo application. *Attributes* are characteristics of the objects in a class: identity, state, configuration parameters, run-time statistics, and so on. There are a number of attributes that are common to MIB operations and replies and common to individual classes. Every class has a *state* attribute that indicates the state of the object. The state of an object is either *return to the user* or *new, changed state*, if you are invoking an operation on the MIB to change an object's state.

Independent of classes is a set of common attributes that are defined in the MIB(5) reference page. These attributes control the input operations, communicate to the MIB what the user is trying to do, and/or identify to the programmer some of the characteristics of the output buffer that are independent of a particular class.

## Using Command-Line Utilities

The BEA Tuxedo system provides a set of commands for managing different parts of the system. The commands enable you to access common administrative utilities. These utilities can be used for the following tasks:

- [Configuring your application using command-line utilities](#)
- [Operating your application using command-line utilities](#)
- [Monitoring your application using command-line utilities](#)

## Configuring Your Application Using Command-line Utilities

You can configure your application by using command-line utilities such as the `vi` text editor. Specifically, you can use command-line utilities to write the configuration file, `UBBCONFIG`, and translate the file from a text format (`UBBCONFIG`) to a binary format (`TUXCONFIG`), by running the `tmloadcf` command. Then you are ready to boot your application.

You can dynamically administer your configuration by adding servers or machines, deleting machines, and so forth. Updating `TUXCONFIG` (the binary file version), however, does not update the `UBBCONFIG` (the text file version). To synchronize both files, you need to back them up. To do this, you translate the binary file back to text by running the `tmunloadcf` command.

**Note:** The `UBBCONFIG` is generated and stored by the application administrator in the application directory (`APPDIR`).

Following is a list of common command-line utilities that you can use to configure your application.

- `tmconfig`—a command that enables you to update some configuration file parameters, or MIB attributes, and add records to some `TUXCONFIG` sections while the BEA Tuxedo system application is running.
- `tmloadcf`—a command that allows you to load the binary `TUXCONFIG` configuration file.
- `tmunloadcf`—a command that allows you to translate the binary configuration file back to a text version, so that `UBBCONFIG` and `TUXCONFIG` can be synchronized.
- `tpacladd`, `tpaclcvt`, `tpacldel`, and `tpaclmod`—a set of commands that allow you to create or manage access control lists for applications. These commands enable the use of security-related authorization features.
- `tpgrpadd`, `tpgrpdel`, `tpgrpmod`—a set of commands that allow you to create and manage user groups by using access control lists to authorize access to services, queues, and events.
- `tpusradd`, `tpusrdel`, `tpusrmod`—a set of commands that allow you to create and manage a user database for authorization purposes.

## See Also

- `UBBCONFIG(5)` in *BEA Tuxedo File Formats and Data Descriptions Reference*
- “Creating the Configuration File” on page 3-22
- “Making Permanent Configuration Changes” on page 3-24

## Operating Your Application Using Command-line Utilities

Once you have configured your application successfully, you can use the following command-line utilities to operate your application.

- `tadmin`—a command that allows you to configure, monitor, and tune a distributed application.
- `tmbot`—a command that allows you to centrally start up your application servers for a distributed application.
- `tmshtdown`—a command that allows you to centrally shut down an application program across a distributed application.

## Managing System Events Using EventBroker

The BEA Tuxedo EventBroker performs the following tasks:

- Monitors events and notifies subscribers when events are posted via `tpost(3c)`.
- Keeps an administrator informed of changes in an application by tracking events.
- Enhances event monitoring by providing a system-wide summary of events.
- Provides a mechanism through which an event can trigger a variety of notification activities.

The EventBroker recognizes over 100 meaningful state transitions in a MIB object as system events. The postings for system events include the current MIB representation of the object on which the event has occurred, and some event-specific fields that identify the event that occurred. For example, if a machine is partitioned, an event is posted with the following information:

- The name of a *machine class object* (`T_MACHINE`), with all the attributes of that machine



- Some event attributes identifying the event as *machine partitioned*

You can use the EventBroker simply by subscribing to system events. Then, instead of having to query for MIB records, you can be informed automatically when events occur in the MIB by receiving FML data buffers representing MIB objects.

## See Also

- “What Is an Event” on page 3-15
- “Subscribing to an Event” on page 3-16
- “Types of Events” on page 3-17
- “Using Event-based Communication” on page 1-14 in *Tutorials for Developing a BEA Tuxedo Application*

## What Is an Event

An event is a state change or other occurrence in a running application that may warrant special attention from an operator, an administrator, or the software. In the EventBroker, events are assigned one of three severity levels:

- *Error*, for example, a server has died or a network connection has been dropped
- *Informational*, for example, a state change has occurred as a result of a process or the detection of a configuration change
- *Warning*, for example, a client has not been allowed to join the application after failing authentication

## Subscribing to an Event


As the administrator for your BEA Tuxedo application, you can enter subscription requests on behalf of a client or server process through calls to the `EVENT_MIB(5)`. You use `tpssubscribe` to subscribe to an event using the EventBroker. You may want to subscribe to Events A, B, and C and request to be informed when they occur.

Each subscription specifies one of the following notification methods:

- **Client notification**—The EventBroker keeps track of the client's interest in these events and a client is notified in the form of unsolicited notification. Some events are anonymously posted. A client can join an application, independent of whether anyone else has subscribed, and post events to the EventBroker. The EventBroker matches these events against its database of subscriptions and sends an unsolicited notification to the appropriate clients.
- **Service calls**—If the subscriber wants event notifications to go to service calls, then the `ctl` parameter must point to a valid `TPEVCTL` structure.
- **Message enqueueing to stable-storage queues**—For subscriptions to stable-storage queues, the queue space, queue name, and correlation identifier are used, in addition to `eventexpr` and `filter`, when determining matches. The correlation identifier can be used to differentiate among several subscriptions for the same event expression and filter rule, destined for the same queue.
- **Placing messages on the ULOG**—Using the `T_EVENT_USERLOG` class of `EVENT_MIB`, subscribers can write system `USERLOG` messages. When events are detected and matched, they are written to the `USERLOG`.
- **Command-line utilities**—Using the `T_EVENT_COMMAND` class of `EVENT_MIB`, the EventBroker tracks and matches events. When a match is found, it is passed to the command used when subscribing to the event.

**Note:** Notification methods are determined by the subscriber process type and the arguments passed to `tpssubscribe`.

Figure 3-3 Subscribing to an Event



## See Also

- `EVENT_MIB(5)` in *BEA Tuxedo File Formats and Data Descriptions Reference*
- `tpssubscribe(3c)` in *BEA Tuxedo C Function Reference*

## Types of Events

The BEA Tuxedo system supports two event types:

- *System Events*—provide details about BEA Tuxedo system events, such as servers dying, and network failures. When an event is posted by clients or servers, the EventBroker matches the posted event's name to a list of subscribers for that event and takes appropriate action, as determined by each subscription.
- *User Events or Application-specific Events*—allow application programs to post events when certain criteria are met. An example is a banking application that posts an event for withdrawals over a certain limit.

## Differences between System and Application-specific Events

The following table identifies the differences between system and application-specific events.

**Table 3-1 Differences between System and Application-specific Events**

Area	Differences
Events	System events are defined in advance by the BEA Tuxedo system code. For an application, designers decide which application events should be monitored. Application programs are written to: (a) detect when an event of interest has occurred, and (b) post the event to the EventBroker through <code>tppost</code> .
Event List	A list of the application event subscriptions is made available to interested users just as the BEA Tuxedo system provides a list of system events available to users with <code>EVENTS(5)</code> . System event names begin with a dot ( <code>.</code> ); application-specific event names may not begin with a dot ( <code>.</code> ).
Subscriptions	Subscribing to an event in an application-specific event broker is similar to subscribing to the BEA Tuxedo System EventBroker. You subscribe by making calls to <code>tpsubscribe</code> using the published list of events for the application. <code>EVENTS(5)</code> lists the notification message generated by an event as well as the event name (used as an argument when <code>tppost</code> is called). Subscribers can use the wild card capability of regular expressions to make a single call to <code>tpsubscribe</code> that covers a whole category of events.

# **BEA Tuxedo Administrative Services**

A set of system servers provides the following administrative services needed by the BEA Tuxedo system:

- [Application queue management](#)
- [Centralized application configuration](#)
- [Distributed application management](#)
- [Dynamic application reconfiguration](#)
- [Event management](#)
- [Security management](#)
- [Startup and shutdown of an application](#)
- [Transaction management](#)
- [Workstation management](#)

## Managing Application Queues

Queueing enables programmers to write applications that communicate by accessing one or more queues. Because of the location transparency of queues, administrators can move queues from one machine to another without requiring any programming changes.

The MIB consists of a queue device, queue spaces, and queues (required by an application), and the BEA Tuxedo system servers that enqueue and dequeue messages from a queue space. Administrators can use the BEA Administration Console or command line utilities to define the queue spaces, queues, and administrative servers in the MIB.

## Using `qmadmin` to Administer Application Queues

The command-line utility `qmadmin` allows you to perform all administration functions for the application queues in a configuration, that is, setting up the universal device list (UDL) and volume table of contents (VTOC) that will contain a queue, defining queue spaces within a queue device, and so on. `qmadmin` enables you to manipulate the file system. Using some run-time monitoring capabilities, you can see how many messages are in queues or how many headers are in messages. You can also change characteristics of queues or messages on queues, delete messages on queues, change the size of devices, and so on. In an application you can have multiple application queue devices, and run application queues on multiple machines. Each machine has its own queue device, so you can run `qmadmin` to monitor and manage a particular application queue device on each machine.

---

Utility	Description
<code>qmadmin</code>	Provides for the creation, inspection and modification of message queues. The name of the device (file) on which the universal device list resides (or will reside) for the queue space may either be specified as a command line argument or through the environment variable <code>QMCONFIG</code> . If both are specified, the command option is used.

---

## Using tmconfig to Modify Your Configuration

The `tmconfig` command enables you to browse and modify the `TUXCONFIG` file and its associated entities, and to add new components (such as machines and servers) while your application is running.

When you modify your configuration file (`TUXCONFIG` on the `MASTER` machine), the `tmconfig` command:

- Updates the `TUXCONFIG` file on all machines in the application that are currently booted.
- Propagates the `TUXCONFIG` file automatically to new machines as they are booted.
- Runs as a BEA Tuxedo system client.

**Note:** Refer to the `tmconfig`, `wtmconfig(1)` and `TM_MIB(5)` in *BEA Tuxedo Command Reference* and *BEA Tuxedo File Formats and Data Descriptions Reference* for information on the semantics, range values, and validation of configuration parameters.

# Managing Your Configuration

The configuration of any application is primarily controlled by the creation and maintenance of a configuration file, or `UBBCONFIG` file. Managing your configuration involves the following tasks:

- [Creating the configuration file to suit your application needs](#)
- [Making permanent configuration changes by updating the `UBBCONFIG` file](#)
- [Changing your configuration while the application is running](#)

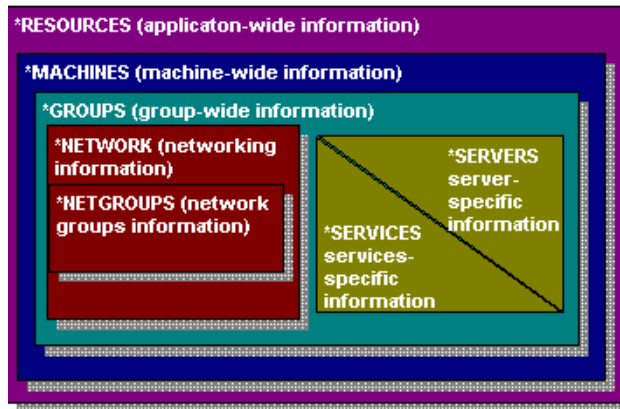
## Creating the Configuration File

Application configuration data is maintained in the `UBBCONFIG`, an ordinary text file on the `MASTER` machine. The configuration file (`UBBCONFIG`) is a repository that contains all the information necessary to boot an application, such as lists of its resources, machines, groups, servers, available services, and so on. Once written, the `UBBCONFIG` file is compiled into a binary file, `TUXCONFIG`. (If you are developing a multi-domain application, you must provide a configuration file for each domain in the application.) An application cannot run without a configuration file.

The `UBBCONFIG` file consists of eight sections, five of which are required for all configurations: `RESOURCES`, `MACHINES`, `GROUPS`, `SERVERS`, and `SERVICES`. The `RESOURCES` and `MACHINES` sections must be the first and second sections, respectively (as illustrated in the following diagram). `GROUPS` must be ahead of `SERVERS` and `SERVICES`.



Figure 3-4 UBBCONFIG File



- RESOURCES—(required) contains system-wide parameters that describe the application as a whole
- MACHINES—(required) contains logical names and types of physical machines
- GROUPS—(required) associates servers with resource managers and machines
- SERVERS—identifies each server in the application
- SERVICES—identifies each service, and specifies priority, loading, and so on
- NETWORK—contains configuration data for LAN environments
- ROUTING—contains data-dependent routing tables
- NETGROUPS—allows for multiple BRIDGES per machine

Your particular configuration determines which sections of the UBBCONFIG file are required. Once you have written your UBBCONFIG file, you must compile it into a binary file called TUXCONFIG. You can generate your TUXCONFIG file by running the `tmloadcf(1)` command or by using the BEA Administration Console.

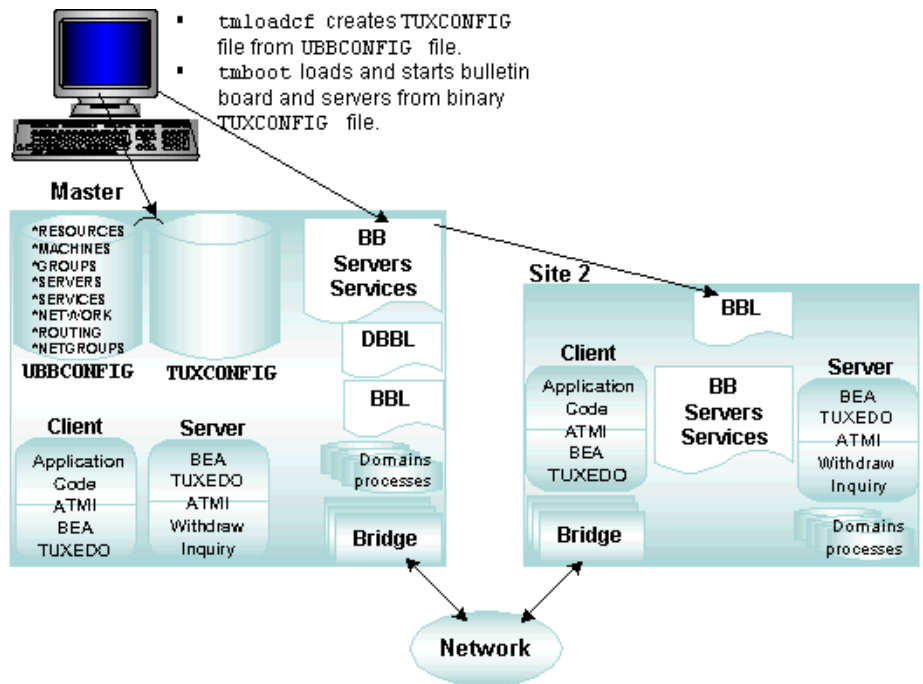
## See Also

- “How to Create a Configuration File” on page 3-2 in *Setting Up a BEA Tuxedo Application*

# Making Permanent Configuration Changes

To make permanent configuration changes, the administrator can use a text editor to update the configuration parameters in the `UBBCONFIG` file, and use the `tmloadcf` utility to load the text file into the binary `TUXCONFIG` file used by the BEA Tuxedo system. When the application is started, `tmboot` loads `TUXCONFIG` into shared memory to establish the bulletin board, propagating the changes to remote machines if necessary.

Figure 3-5 Configuration Management



# Managing Your Configuration Dynamically

Administrators can use the BEA Administration Console or the BEA Tuxedo system command-line utilities to reconfigure applications dynamically, adjusting parameters to respond to varying system loads while the system is running. A revised `TUXCONFIG` file is propagated automatically to all machines in the system as it is updated. However, many `RESOURCES` parameters cannot be changed while the system is running.

Examples of tasks you can do dynamically include: adding servers or machines, deleting machines, and so forth. To ensure that the text and binary versions of your configuration file (`UBBCONFIG` and `TUXCONFIG`, respectively) always match, you need to back them up and synchronize them by using `tmunloadcf`. This command translates the binary file to a text version.

You can change most elements of the system dynamically. You can, for example, spawn new servers, add new machines, or change time-out parameters. There are, however, a few things you cannot change while a system is running:

- Any parameter in the configuration file that affects the size and shape of the bulletin board cannot be changed. Many such parameters are named with the prefix “MAX,” such as the `MAXGTT` parameter, which specifies the maximum number of in-flight transactions allowed within the BEA Tuxedo system at any time.
- The processor name of a machine within a particular application cannot be changed. (You can add new machines with different names but you cannot change the name of an existing machine.)
- The values of server executables, assigned to run on `MASTER` and `BACKUP` machines, cannot be changed.

## See Also

- “Performing Dynamic Operations Using `tmadmin(1)`” on page 3-26

# Performing Dynamic Operations Using `tmadmin(1)`

Using the `tmadmin(1)` command, you can perform any of the following operations to a running application:

- Monitor performance by checking statistics on groups, servers, and services (`bbstats`, `bbparms`)
- Modify server and service parameters such as those that change load values (`changeload`), suspend and resume services (`suspend` and `resume`), advertise and unadvertise services (`advertise` and `unadvertise`), and change the AUTOTRAN time-out value (`changetrancode`)
- Boot (`boot`), cleanup (`pclean`), and migration (`migratemach`, `migrategroup`)

## Commonly Used `tmadmin` Commands

`tmadmin` provides subcommands that enable you to monitor your run-time system, tune your application, and dynamically configure your application. Following is a list of the most commonly used `tmadmin` commands. (For a comprehensive list of the `tmadmin` commands, refer to the `tmadmin(1)` in *BEA Tuxedo Command Reference*.)

- `help` - provides you with a list of subcommands, their abbreviation, arguments, and descriptions.
- `printserver` (`psr`)—prints information for application and administrative servers.
- `printservice` (`psc`)—prints information for application and administrative services.
- `printclient` (`pclt`)— prints information for the specified set of client processes. If no arguments or defaults are set, then information on all clients is printed.

## Sample Output from the tadmin Command

Following is sample output from the tadmin printserver (psr) command, which provides information about application and administrative servers.

**Figure 3-6 Sample Output from the tadmin printserver Command**

```
>psr
```

Prog Name	Queue Name	Grp Name	ID	RqDone	Load Done	Current Service
BBL	83108	SITE1	0	1	50	( IDLE )
AUDITC	auditc	BANKB1	1	0	0	( IDLE )
XFER	00001.00101	BANKB1	101	1	30	( TRANSFER )
TMS_SQL	BANKB1_TMS	BANKB1	30001	0	0	( IDLE )
ACCT	00001.00102	BANKB1	102	0	0	( IDLE )
TMS_SQL	BANKB1_TMS	BANKB1	30002	0	0	( IDLE )
BAL	00001.00103	BANKB1	103	6	7	( IDLE )
BTADD	00001.00104	BANKB1	104	0	0	( IDLE )
BALC	00001.00105	BANKB1	105	0	0	( IDLE )
TLR	tlr1	BANKB1	111	0	0	( IDLE )
TLR	tlr1	BANKB1	112	3	110	( WITHDRAWAL )
TLR	tlr1	BANKB1	113	0	0	( IDLE )
TLR	tlr1	BANKB1	114	0	0	( IDLE )
TLR	tlr1	BANKB1	115	0	0	( IDLE )
TLR	tlr1	BANKB1	116	9	100	( IDLE )
TLR	tlr1	BANKB1	117	20	2048	( IDLE )
TLR	tlr1	BANKB1	118	30	600	( IDLE )
TLR	tlr1	BANKB1	119	0	0	( IDLE )
TLR	tlr1	BANKB1	120	0	0	( IDLE )

```
>
```

## See Also

- “How a tadmin Session Works” on page 2-13 in *Administering a BEA Tuxedo Application at Run Time*
- “Using Command-line Utilities to Monitor Your Application” on page 2-10 in *Administering a BEA Tuxedo Application at Run Time*

# Managing a Distributed Application Centrally

Even if your BEA Tuxedo application is large and complex, you can perform all run-time administrative functions from one `MASTER` machine. You can do so using the BEA Tuxedo system-supplied command-line utilities, or the BEA Administration Console, or through your third-party administration tools used with the BEA Manager product.

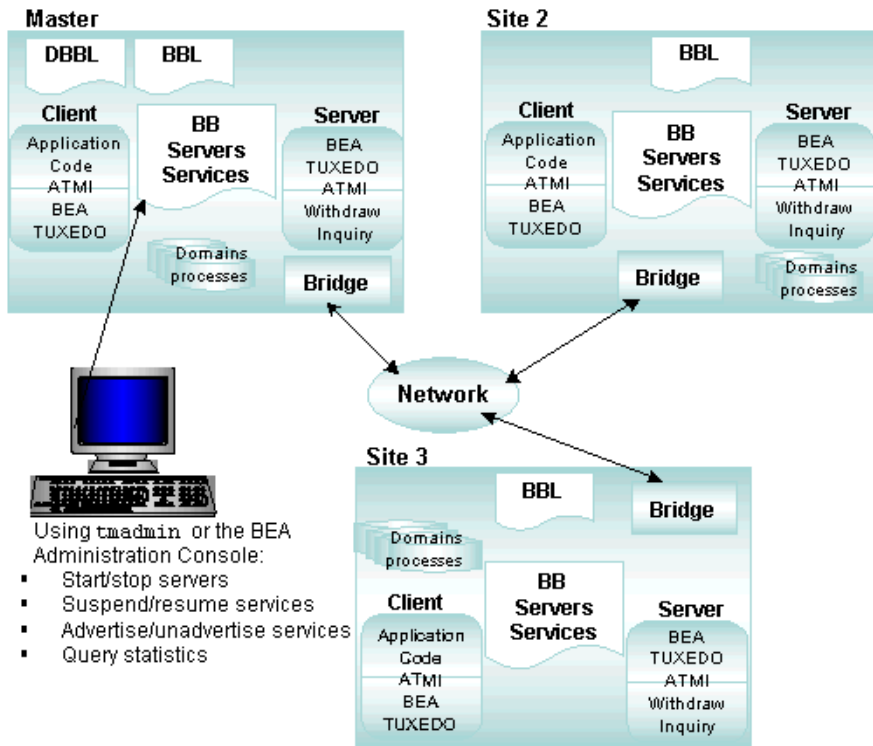
From the `MASTER` machine, you can configure your application, initiate start-up and shutdown, and perform administrative tasks during runtime. All other machines can query the `MASTER` machine. From the `MASTER` machine, you have control over configuration, fault management, security, monitoring, and performance.

You can use the following two methods to make changes to your system while it is running:

- [The BEA Administration Console](#)—a graphical user interface (GUI) to the commands that perform administrative tasks, including dynamic system modification
- [The `tmadmin` command](#)—a shell-level meta-command that enables you to run 50 subcommands for performing various administrative tasks, including dynamic system modification

Because it is a graphical user interface, the BEA Administration Console is simpler to use than the `tmadmin` command interpreter. If you prefer using a GUI, bring the BEA Administration Console up on your screen as soon as you are ready to begin an administrative task. Graphics and online help provided with the BEA Administration Console guide you through any task you need to perform. The following illustration shows how you can use the `tmadmin` command or the BEA Administration Console to control a run-time application. All operations can be performed from the `MASTER` machine. The utilities directly affect the bulletin board on the `MASTER` machine, and updates are distributed to other bulletin boards automatically.

Figure 3-7 Centralized Control of a Distributed Application



## See Also

- “Using the BEA Administration Console” on page 3-4
- “Performing Dynamic Operations Using `tmadmin(1)`” on page 3-26

## Managing Security

Administrators can configure applications with appropriate levels of security provided by the BEA Tuxedo system. Incremental levels of authentication and authorization can be used to define access to an application. Levels can vary from *no authentication* for highly secure environments, to a password or an access control list (ACL) that filters who can use services, post an event, and enqueue or dequeue a message on a queue.

With an ACL, not only is a user authenticated when joining an application, but permissions are checked automatically when attempts are made to access application entities, such as services. When an ACL is created for a resource, users not included on the list are denied access to the resource. Resources unprotected by an ACL are accessible by any client who successfully joins the application. Resources unprotected by an ACL with the `MANDATORY_ACL` security option specified, are denied for any client who joins the application.

An application can be configured so that all servers (except `AUTHSVR`, the BEA Tuxedo administration server) have restricted access to shared resources, such as shared memory and message queues. When a client joins an application, `AUTHSVR` provides an authentication service that verifies whether the user has the correct authentication level (in the MIB). This service is transparent to the programmer.

## See Also

- “Selecting Security Options” on page 3-31
- “Setting Up Security” on page 3-32
- “Administering Security” on page 2-1 in *Using BEA Tuxedo Security*
- “Programming Security” on page 3-1 in *Using BEA Tuxedo Security*



---

# Selecting Security Options

The following are the security options provided by the BEA Tuxedo system:

- *No authentication*—Clients do not have to be verified before joining an application.
- *Application Password*—A single password is defined for an entire application and clients must provide the password to join the application.
- *User-level Authentication*—In addition to an application password, each client must provide a valid user name and application-specific data such as a password to join the application.
- *Optional Access Control List (ACL)*—Clients must provide an application password, a user name, and a user password. If there is no ACL associated with a user name, permission is granted. This practice enables an administrator to configure access for only those resources that need more security; ACLs need not be configured for services, queues, or events that are open to everyone.
- *Mandatory Access Control List (ACL)*—Clients must provide an application password, a user name, and a user password. This level is similar to optional ACL, but an access control list must be configured for every entity (such as a service, queue, or event) that users can access. If mandatory ACLs are being used and there is no ACL for a particular entity, permission for that entity is denied.
- *Link-Level Encryption*—Users of BEA Tuxedo System Security can establish data privacy for messages moving over the network links that connect the machines in a BEA Tuxedo application. The BEA Tuxedo system encrypts data before sending it over a network link and decrypts it as it comes off the link. Three levels of security are offered: 0-bit (no encryption), 56-bit (international), or 128-bit (US and Canada).
- *Public key encryption*—Consists of message-based encryption and message-based digital signature. Message-based encryption reveals user data only to designated recipients. With message-based digital signature, a sending process must prove its identity, and bind that proof to a specific message buffer. Any third party can verify the signature's authenticity. Undetected tampering is impossible because a digital signature contains a cryptographically secure

checksum computed on the entire contents of a buffer. A digital signature also contains a tamper-proof stamp based on the originating machine's local clock.

- *Auditing*—Collects, stores, and distributes information about operating requests and their outcomes.

## Setting Up Security

The type of administrative work and/or programming you must do to set up security for your application depends upon the security options that you choose. Administratively, you need to configure the MIBs using either the BEA Administration Console or the command-line utilities.

You can also build your own security mechanisms. To do so, set the application security level to *User-Level Authentication* and specify an application service that performs authentication in the BEA Tuxedo MIB.

To enable authentication and authorization, administrators must configure the following in the MIB:

- AUTHSVR server
- Identity and passwords of authorized users
- Access control lists used on services, queues, and/or events

# Starting Up and Shutting Down your Application

To *start an application*, you need to perform the following tasks as stated in *Administering a BEA Tuxedo Application at Run Time*.

1. Set the environment variables as described in “How to Set Your Environment” on page 1-2
2. Create the `tuxconfig` file as described in “How to Create the TUXCONFIG File” on page 1-4
3. Propagate the BEA Tuxedo software as described in “How to Propagate the BEA Tuxedo System Software” on page 1-5
4. Create a `TLOG` device, (if required) as described in “How to Create a TLOG Device” on page 1-6
5. Start `tlisten` at all sites (MP environments) as described in “How to Start tlisten at All Sites” on page 1-7
6. Boot the application as described in “How to Boot the Application” on page 1-9

To *shut down an application*, you need to perform the following task.


Run `tmshutdown` on the `MASTER` machine as described in “How to Shut Down Your Application” on page 1-11

# Managing Transactions

A powerful feature of the BEA Tuxedo system is the ability to manage transactions for database applications that support the XA-interface. Transactions simplify the writing of distributed applications. They allow your application to cope more easily with a large set of problems that can occur in distributed environments, such as machine, program, and network failures.

In a distributed architecture, a local machine involved in a transaction can communicate with a remote machine which may, in turn, communicate with another remote machine. The communication and the work done by the remote machines is part of the transaction, and integrity must be maintained. Keeping track of distributed transaction processing (DTP) can be a complex task because the system must maintain enough information about a transaction to be able to roll it back (that is, to undo it) at any moment.

**Figure 3-8 Transaction Management**



To keep track of the participants in a transaction, the BEA Tuxedo system creates a transaction log. To maintain the state of an application as represented by the contents of the computer's memory, the BEA Tuxedo system uses one or more Resource Managers (a collection of information and processes for accessing it, such as a database management system). To coordinate all the operations performed and all the modules affected by a transaction, the BEA Tuxedo system uses a Transaction Manager (TM), which directs the actions of the RMs. Together, TMs and RMs maintain the atomicity of a distributed transaction.

## Coordinating Operations with a Transaction Manager Server (TMS)

The BEA Tuxedo Transaction Manager (TM) is responsible for coordinating global transactions involving system-wide resources. Local resource managers (RMs) are responsible for individual resources. The transaction manager server (TMS) begins, commits, and aborts transactions involving multiple resources. The server uses an embedded SQL interface to the RM to read and update the database accessed by the server group. The TMS and RMs use the XA interface to perform all or none of the resource work in a *global transaction*.

## Tracking Participants with a Transaction Log (TLOG)

A global transaction is logged in the transaction log (TLOG) only when it is in the process of being committed. The TLOG records the reply from the global transaction participants at the end of the first phase of a 2-phase-commit protocol. A TLOG record indicates that a global transaction should be committed; no TLOG record is written for those transactions that are to be rolled back. In the first phase, or pre-commit, each Resource Manager must commit to performing the transaction request. Once all parties commit, transaction management commits and completes the transaction. If either task fails because of an application or system failure, both tasks fail and the work performed is undone or “rolled back” to its initial state.

The TMS that coordinates global transactions uses the TLOG file. Each machine should have its own TLOG.

**Note:** Customers using the Domains feature should note that the Domains gateway performs the functions of the TMS in Domains groups. However, Domains uses its own transaction log containing information similar to that in the TLOG, in addition to Domains-specific information.

## See Also

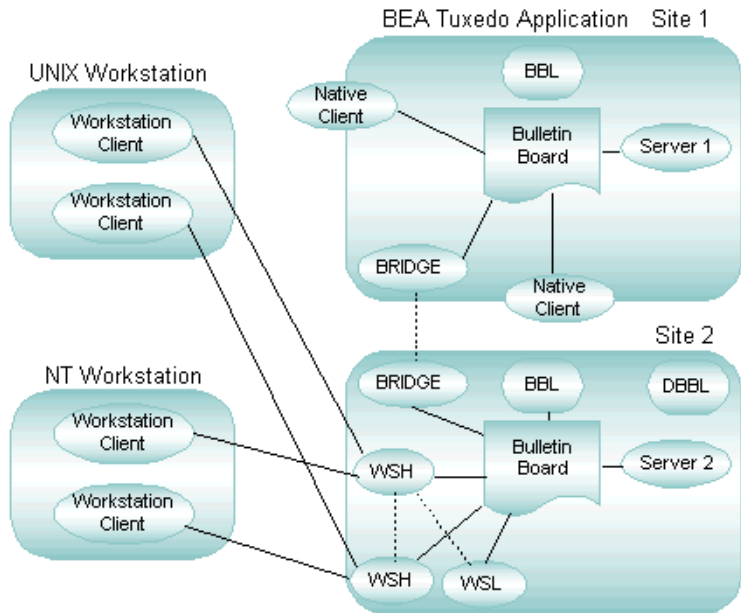
- “Using Transactions” on page 1-18 in *Tutorials for Developing a BEA Tuxedo Application*
- “Configuring Your Application to Use Transactions” on page 5-1 in *Setting Up a BEA Tuxedo Application*

# Managing Workstations

Workstation clients need enough of the BEA Tuxedo system software to package the information associated with a request. They can then send that information to a system that supports all the BEA Tuxedo system software, including ATMI functions and networking software.

The administrator configures one or more Workstation Listeners (WSLs) to be ready for connection requests from Workstation clients. Each WSL uses one or more associated Workstation Handlers (WSHs) to handle the client's workload. Each WSH manages multiple workstations, multiplexes all communication with a particular workstation over a single connection.

**Figure 3-9 Handling Workstation Clients**



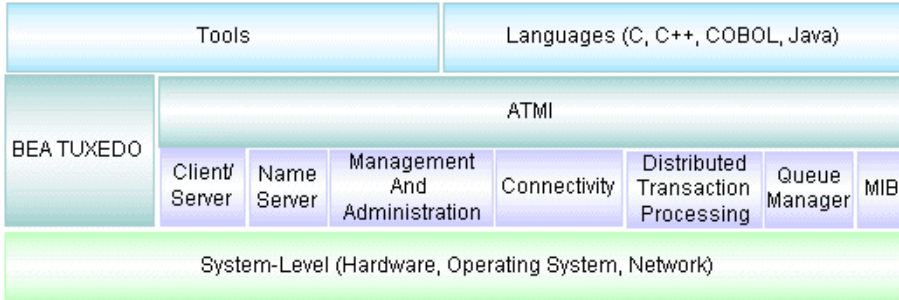
A machine can then handle thousands of workstation clients. An administrator can define several WSLs in a domain to distribute and balance the workstation communication load across multiple machines.

From a programming perspective, all client ATMI programming interfaces are supported for workstation client development.

## Development View: What You Can Do Using the ATMI

The Application to Transaction Monitor Interface (ATMI), the BEA Tuxedo API, provides an interface for communications, transactions, and management of data buffers that works in all environments supported by the BEA Tuxedo system. It defines the interface between application programs and the BEA Tuxedo system. The ATMI offers a simple interface for a comprehensive set of capabilities. It implements the X/Open DTP model of transaction processing.

**Figure 3-10 Using the ATMI**



The ATMI supports the following tasks:

- Client initialization
- Server naming
- System messaging
- Managing transactions
- Dispatching of services
- Managing buffers

The ATMI library offers you a variety of functions for defining and controlling global transactions in a BEA Tuxedo application. Global transactions enable you to manage exclusive units of work spanning multiple programs and resource managers in a distributed application. All the work in a single transaction is treated as a logical unit,

### 3 Three Ways of Viewing the BEA Tuxedo System Infrastructure

---

so that if any one program cannot complete its task successfully, no work is performed by any programs in the transaction. Most ATMI functions support different communication styles. These functions knit together distributed programs by enabling them to exchange data. All ATMI functions send or receive data in typed buffers. Following is a list of ATMI functions (for C and COBOL bindings), grouped by the type of task they perform.

**Table 3-2 Using the ATMI Functions**

<b>For a Task Related to</b>	<b>Use This C Function</b>	<b>Or This COBOL Function</b>	<b>To</b>
Client membership	<code>tpchkauth(3c)</code>	<code>TPCHKAUTH(3cbl)</code>	Check whether authentication is required
	<code>tpinit(3c)</code>	<code>TPINITIALIZE(3cbl)</code>	Have a client join an application
	<code>tpterm(3c)</code>	<code>TPTERM(3cbl)</code>	Have a client leave an application
Buffer management	<code>tpalloc(3c)</code>	N/A	Create a message buffer
	<code>tprealloc(3c)</code>	N/A	Resize a message buffer
	<code>tpfree(3c)</code>	N/A	Free a message buffer
	<code>tpypes(3c)</code>	N/A	Get a message type and subtype
Message priority	<code>tpgprio(3c)</code>	<code>TPGPRIOR(3cbl)</code>	Get the priority of the last request
	<code>tpsprio(3c)</code>	<code>TPSPRIOR(3cbl)</code>	Set the priority of the next request
Request/Response communications	<code>tpcall(3c)</code>	<code>TPCALL(3cbl)</code>	Initiate a synchronous request/response to a service
	<code>tpacall(3c)</code>	<code>TPACALL(3cbl)</code>	Initiate an asynchronous request (fanout)
	<code>tpgetrply(3c)</code>	<code>TPGETRPLY(3cbl)</code>	Receive an asynchronous response
	<code>tpcancel(3c)</code>	<code>TPCANCEL(3cbl)</code>	Cancel an asynchronous request



## Development View: What You Can Do Using the ATMI

For a Task Related to	Use This C Function	Or This COBOL Function	To
Conversational communications	<code>tpconnect(3c)</code>	<code>TPCONNECT(3cbl)</code>	Begin a conversation with a service
	<code>tpdiscon(3c)</code>	<code>TPDISCON(3cbl)</code>	Abnormally terminate a conversation
	<code>tpsend(3c)</code>	<code>TPSEND(3cbl)</code>	Send a message in a conversation
	<code>tprecv(3c)</code>	<code>TPRECV(3cbl)</code>	Receive a message in a conversation
Reliable queuing	<code>tpenqueue(3c)</code>	<code>TPENQUEUE(3cbl)</code>	Enqueue a message to a message queue
	<code>tpdequeue(3c)</code>	<code>TPDEQUEUE(3cbl)</code>	Dequeue a message from a message queue
Event-based communications	<code>tpnotify(3c)</code>	<code>TPNOTIFY(3cbl)</code>	Send an unsolicited message to a client
	<code>tpbroadcast(3c)</code>	<code>TPBROADCAST(3cbl)</code>	Send messages to several clients
	<code>tpsetunsol(3c)</code>	<code>TPSETUNSOL(3cbl)</code>	Set unsolicited message call-back
	<code>tpchkunsol(3c)</code>	<code>TPCHKUNSOL(3cbl)</code>	Check the arrival of unsolicited messages
	N/A	<code>TPGETUNSOL(3cbl)</code>	Get an unsolicited message
	<code>tppost(3c)</code>	<code>TPPOST(3cbl)</code>	Post an event message
	<code>tpsubscribe(3c)</code>	<code>TPSUBSCRIBE(3cbl)</code>	Subscribe to event messages
	<code>tpunsubscribe(3c)</code>	<code>TPUNSUBSCRIBE(3cbl)</code>	Unsubscribe to event messages

### 3 Three Ways of Viewing the BEA Tuxedo System Infrastructure

---

<b>For a Task Related to</b>	<b>Use This C Function</b>	<b>Or This COBOL Function</b>	<b>To</b>
Transaction management	<code>tpbegin(3c)</code>	<code>TPBEGIN(3cbl)</code>	Begin a transaction
	<code>tpcommit(3c)</code>	<code>TPCOMMIT(3cbl)</code>	Commit the current transaction
	<code>tpabort(3c)</code>	<code>TPABORT(3cbl)</code>	Roll back the current transaction
	<code>tpgetlev(3c)</code>	<code>TPGETLEV(3cbl)</code>	Check whether in transaction mode
	<code>tpsuspend(3c)</code>	<code>TPSUSPEND(3cbl)</code>	Suspend the current transaction
	<code>tpresume(3c)</code>	<code>TPRESUME(3cbl)</code>	Resume a transaction
	Service entry and return	<code>tpsvrinit(3c)</code>	<code>TPSVRINIT(3cbl)</code>
<code>tpsvrdone(3c)</code>		<code>TPSVRDONE(3cbl)</code>	Terminate a server
<code>tpservice(3c)</code>		N/A	Prototype for a service entry point
N/A		<code>TPSVCSTART(3cbl)</code>	Get service information
<code>tpreturn(3c)</code>		<code>TPRETURN(3cbl)</code>	End a service function
<code>tpforward(3c)</code>		<code>TPFORWAR(3cbl)</code>	Forward request
Dynamic advertisement		<code>tpadvertise(3c)</code>	<code>TPADVERTISE(3cbl)</code>
	<code>tpunadvertise(3c)</code>	<code>TPUNADVERTISE(3cbl)</code>	Unadvertise a service name
Resource management	<code>tpopen(3c)</code>	<code>TPOPEN(3cbl)</code>	Open a resource manager
	<code>tpclose(3c)</code>	<code>TPCLOSE(3cbl)</code>	Close a resource manager

**Note:** The use of ATMI transaction management functions is optional.

## See Also

- “Using the ATMI to Handle System and Application Errors” on page 2-28 in *Administering a BEA Tuxedo Application at Run Time*
- “Creating a BEA Tuxedo Client” on page 1-2 in *Tutorials for Developing a BEA Tuxedo Application*
- “Creating a BEA Tuxedo Server” on page 1-4 in *Tutorials for Developing a BEA Tuxedo Application*
- “Using Typed Buffers in Your Application” on page 1-6 in *Tutorials for Developing a BEA Tuxedo Application*
- “What Are the BEA Tuxedo Messaging Paradigms” on page 2-8
- “What Is Meant by Naming” on page 2-42

# Run-Time System View: Using Tools in Different Configurations

The BEA Tuxedo system provides tools to create, monitor, and manage both processes and the communication that occurs between processes in a given application. You can use the basic processes and messaging paradigms in many different configurations. Each configuration falls into one of the following run-time categories:

- **Single machine application**—One or more local or remote clients communicate with one or more servers residing on the same machine.
- **Distributed application across multiple machines**—One or more local or remote clients communicate with one or more servers residing on several machines in one domain.
- **Multiple-domain application**—Two or more domains communicate with each other.

## Run-Time System Capabilities

The following table lists the BEA Tuxedo system functionality available in a single-machine application, a distributed application, and a multiple-domain application.

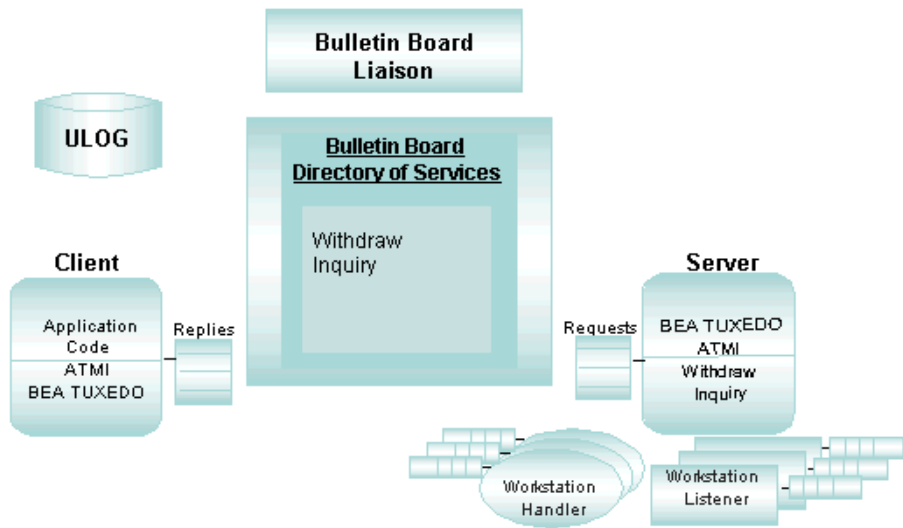
**Table 3-3 Functionality Available in Different Types of Configurations**

Available Functionality	Single-machine Configuration	Multiple-machine (Distributed) Configuration	Multiple-domain Configuration
ATMI	X	X	X
Messaging paradigms	X	X	X
<i>Administrative Parts:</i>			
Bulletin Board (BB), Bulletin Board Liaison (BBL), TLOG, UBBCONFIG, ULOG, TUXCONFIG	X	X	X
Distinguished Bulletin Board Liaison (DBBL)		X	X
Bridges		X	X
<i>Domains processes:</i>			
DMADM, GWADM, GWTDOMAIN (for TDomains), dmloadcf, dmunloadcf, and DMCONFIG, DMTLOG and BDMCONFIG			X
<i>Application processes: clients, servers, and services</i>			
Application processes: <i>clients, servers, and services</i>	X	X	X
Queuing	X	X	X
Transaction management	X	X	X
Event management	X	X	
Security management	X	X	X

# What Is a Single-machine Configuration

A single-machine configuration consists of one or more local or remote clients that communicate with one or more servers residing on a single machine running one or more business applications. Even though it may include multiple applications, this type of configuration is considered a single domain because it is administered as a single entity. All the managed elements (services, servers, and so on) of all the applications in this configuration are defined in and controlled from one BEA Tuxedo configuration file. The basic parts of a single-machine configuration when installed and running on a single machine are illustrated in the following diagram.

**Figure 3-11 A Single-machine BEA Tuxedo Configuration**



### 3 Three Ways of Viewing the BEA Tuxedo System Infrastructure

---

**Table 3-4 Parts of a Single-machine Configuration**

<b>Single-machine Part</b>	<b>Description</b>
Bulletin Board (BB)	A shared memory segment that holds configuration and dynamic information for the system. It is available to all BEA Tuxedo processes.
Bulletin Board Liaison (BBL)	A BEA Tuxedo administrative process that monitors both the data stored in the bulletin board (including any changes made to it) and all application programs.
Clients	Executable programs that periodically request services through the BEA Tuxedo system. (Client programs are normally written by customers.)
Message queues	Communication between clients and servers is performed through operating-system supported, memory-based message queues.
Messaging paradigms	Different models of transferring messages between a client and a server. Examples include request/response mode, conversational mode, events, and unsolicited communication.
Servers	Executable programs that offer named services through the BEA Tuxedo system. (Server programs are normally written by customers.)
Workstation Handler (WSH)	A multi-contexted gateway process on a server that manages service requests from Workstation clients (that is, client processes running on remote sites).
Workstation Listener (WSL)	A server process running on an application site that listens for and distributes connections from Workstation clients (client process running on a remote site).
ULOG (User Log)	A file in which error messages are stored.

## See Also

- “How to Create a Configuration File” on page 3-2 in *Setting Up a BEA Tuxedo Application*

# What Is a Multiple-machine (Distributed) Configuration

A distributed-domain (or multiple-machine) configuration consists of one or more business applications running on multiple machines. Although it includes multiple machines, this type of configuration is considered a single domain because it is administered centrally as a single entity. In other words, all the elements (services, servers, machines, and so on) of all the applications on all the machines in this configuration are defined in, and controlled from, one BEA Tuxedo configuration file.

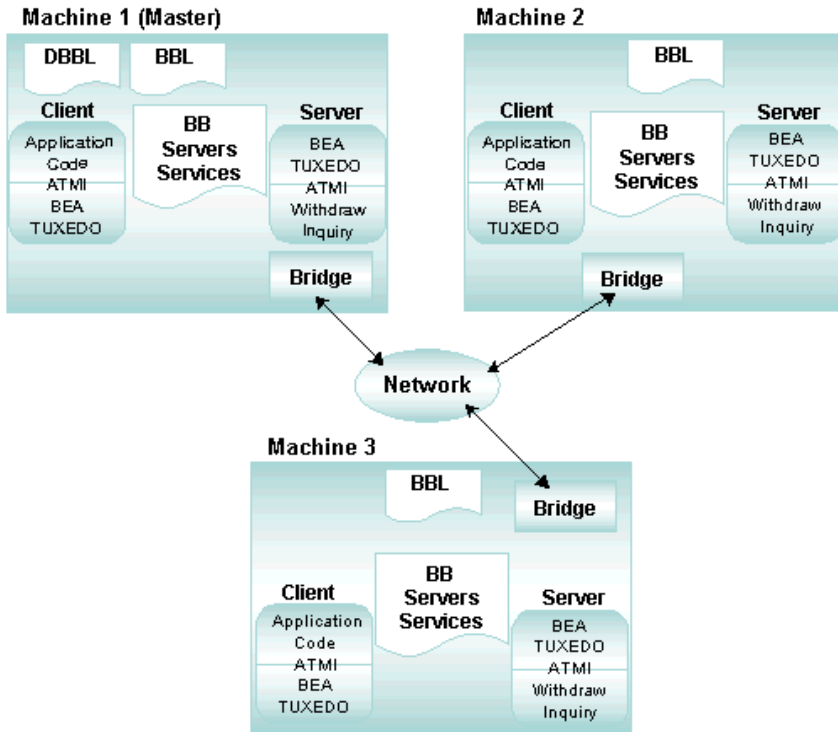
As a business grows, application developers may need to organize different segments of the business by sets of functionality that require administrative autonomy but allow sharing of services and data. Each functionality set defines an application that may span one or more machines, and that is administered independently from other applications. Such a functionally distinct application is referred to as a domain.

The names of domains frequently reflect the functionality provided. When domains have names such as “marketing” and “research and development,” it is easy for customers to find the applications they need.

### 3 Three Ways of Viewing the BEA Tuxedo System Infrastructure

The basic parts of a configuration distributed across multiple machines are illustrated in the following diagram.

**Figure 3-12 Distributed Application**



**Table 3-5 Parts of a Distributed Configuration**

Multiple Machine Part	Description
<a href="#">BRIDGES</a>	BEA Tuxedo system-supplied servers within a domain that send and receive service requests between machines, and route requests to local servers (literally, to local server queues).



## What Is a Multiple-machine (Distributed) Configuration

---

<b>Multiple Machine Part</b>	<b>Description</b>
Bulletin Board (BB)	A shared memory segment that holds configuration and dynamic information for the system. It is available to all BEA Tuxedo processes.
Bulletin Board Liaison (BBL)	A BEA Tuxedo administrative process that monitors both the data stored in the bulletin board (including any changes made to it), and all application programs.
Clients	Executable programs that periodically request services through the BEA Tuxedo system. (Client programs are usually by customers.)
Distinguished Bulletin Board Liaison (DBBL)	A process dedicated to making sure that the BBL server on each machine is alive and functioning correctly. This server runs on the <i>Master</i> machine of a domain and communicates directly with all administration facilities.
Message queues	Communication between clients and servers is performed through operating-system supported, memory-based message queues.
Messaging paradigms	Different models of transferring messages between a client and a server. Examples include request/response mode, conversational mode, events, and unsolicited communication.
Servers	Executable programs that offer named services through the BEA Tuxedo system. (Server programs are normally written by customers.)
Workstation Handler (WSH)	A multi-contexted gateway process on a server that manages service requests from Workstation clients (that is, client processes running on remote sites).
Workstation Listener (WSL)	A server process running on an application site that listens for and distributes connections from Workstation clients (client processes running on remote sites).
ULOG (User Log)	A file in which error messages are stored.

### 3 *Three Ways of Viewing the BEA Tuxedo System Infrastructure*

---

A configuration that runs on more than one machine requires platform interoperability and server transparency.

- *Platform interoperability* means that your application can rely on intermachine communications even when different machines are running different operating systems, without code customization.
- *Server transparency* means that a client can access a server without specifying its location. The locations of servers are recorded in the bulletin board and accessed as needed. As a result, servers can be moved, dropped, or added to an application dynamically, without needing to change the application itself.

The DBBL and BRIDGE servers support these requirements of a distributed-domain configuration.

## See Also

- “How to Create the Configuration File for a Multiple-machine (Distributed) Application” on page 3-3 in *Setting Up a BEA Tuxedo Application*
- “Distributing Applications Across a Network” on page 6-1 in *Setting Up a BEA Tuxedo Application*
- “Creating the Configuration File for a Distributed Application” on page 7-1 in *Setting Up a BEA Tuxedo Application*
- “Setting Up the Network for a Distributed Application” on page 8-1 in *Setting Up a BEA Tuxedo Application*
- “Managing the Network in a Distributed Application” on page 4-1 in *Administering a BEA Tuxedo Application at Run Time*

# What Is a Multiple-domain Configuration

A multiple-domain configuration consists of two or more domains that communicate with each other. Each domain may be either a single-machine configuration or a multiple-machine configuration. Inter-domain communication is achieved through a highly asynchronous multitasking gateway that processes outgoing and incoming service requests to or from all domains. Multiple BEA Tuxedo domains can be connected, allowing clients in one domain transparent access to services physically located in remote domains. Each domain can share services and data, but is administered separately.

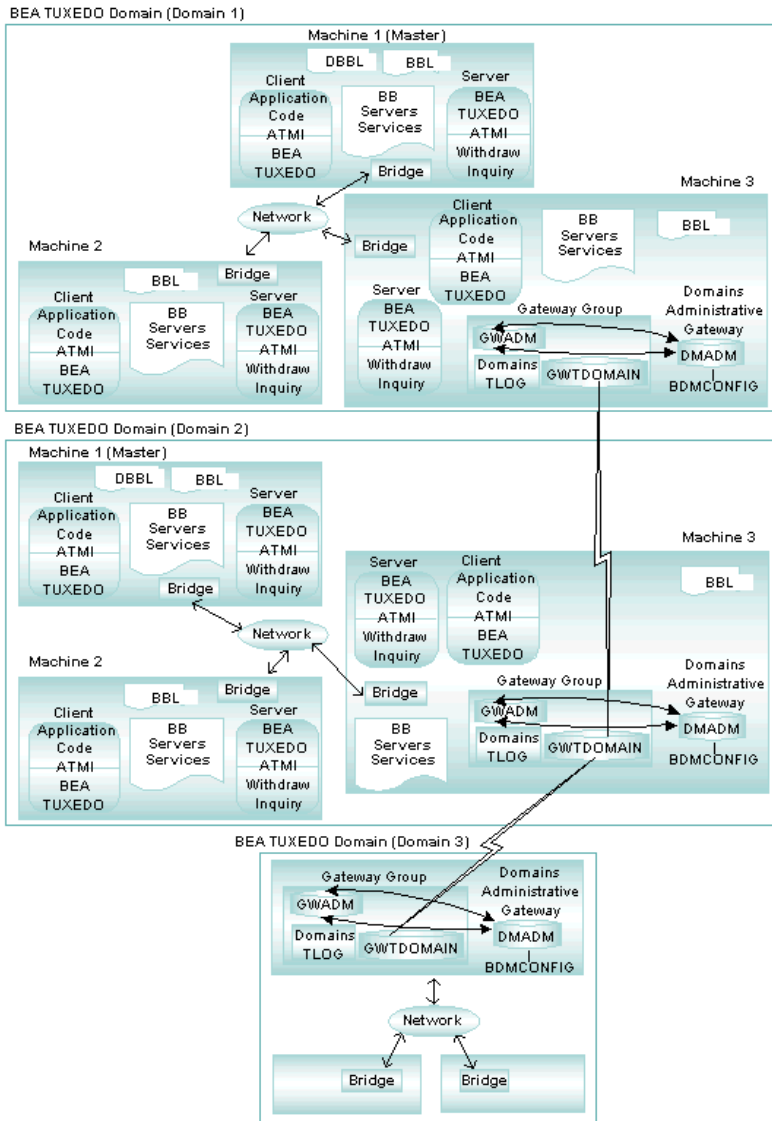
The BEA Tuxedo system provides different types of gateways to accommodate various network transport protocols. Following are the different types of Domains gateways:

- The BEA Tuxedo Domains (TDomains) gateway provides interoperability between two or more BEA Tuxedo applications through a specially designed TP protocol that flows over network transport protocols such as TCP/IP.
- The BEA eLink for Mainframe-OSI TP gateways provides interoperability between BEA Tuxedo applications and other transaction processing applications that use the OSI TP standard. OSI TP is a protocol for distributed transaction processing defined by the International Standards Organization (ISO).
- The BEA eLink for Mainframe-SNA gateway provides interoperability between clients and servers in a BEA Tuxedo domain and clients and servers in an MVS/CICS or MVS/IMS environment in remote SNA domains. It also connects a local BEA Tuxedo domain to multiple SNA networks.
- The BEA eLink for Mainframe-TCP for CICS is a gateway connectivity feature that makes it possible for non-transactional tasks within BEA Tuxedo regions to access services provided by CICS application programs and vice-versa. It enables a BEA Tuxedo domain to communicate via the TCP/IP network transport protocol to a CICS environment.  
The BEA eLink for Mainframe-TCP for IMS is a gateway connectivity feature that provides transparent communications between client and server transactions in an IMS system and a BEA Tuxedo domain, a CICS system, or another IMS system.
- The TOP END Domain Gateway (TEDG) provides interoperability between BEA TOP END systems and BEA Tuxedo domains.

### 3 Three Ways of Viewing the BEA Tuxedo System Infrastructure

The basic parts of a multiple-domain configuration are illustrated in the following diagram.

**Figure 3-13 Multiple-domain Configuration**



**Table 3-6 Parts of a Multiple-domain Configuration**

Multiple-domain Part	Description
BRIDGES	BEA Tuxedo system-supplied servers within a domain that send and receive service requests between machines, and route requests to local servers (literally, to local server queues).
Bulletin Board (BB)	A shared memory segment that holds configuration and dynamic information for the system. It is available to all BEA Tuxedo processes.
Bulletin Board Liaison (BBL)	A BEA Tuxedo administrative process that monitors both the data stored in the bulletin board (including any changes made to it), and all application programs.
Clients	Executable programs that periodically request services through the BEA Tuxedo system. (Client programs are normally written by customers.)
Distinguished Bulletin Board Liaison (DBBL)	Ensures that the BBL servers on each machine are alive and functioning correctly. This server runs on the <i>Master</i> machine of an application and communicates directly with any administration facility.
Domains tools: DMADM, GWADM, GWTDOMAIN, dmloadcf, dmunloadcf, and DMCONFIG	<ul style="list-style-type: none"> <li>■ DMADM-the Domains administrative server</li> <li>■ GWADM-the gateway group administrative server that registers with the DMADM server to obtain configuration information used by the gateway group</li> <li>■ GWTDOMAIN-the gateway process that provides connectivity to remote gateway processes (for TDomains)</li> <li>■ dmloadcf -translates the DMCONFIG file to a binary BDMCONFIG configuration file</li> <li>■ dmunloadcf-translates the BDMCONFIG configuration file from the binary representation into ASCII</li> <li>■ DMCONFIG-the Domains configuration file</li> </ul>
Message queues	Communication between clients and servers is performed through operating-system supported, memory-based message queues.

### 3 Three Ways of Viewing the BEA Tuxedo System Infrastructure

---

Multiple-domain Part	Description
<a href="#">Messaging paradigms</a>	Different models of transferring messages between a client and a server. Examples include request/response mode, conversational mode, events, and unsolicited communication.
<a href="#">Servers</a>	Executable programs that offer named services through the BEA Tuxedo system. (Server programs are normally written by customers.)
<a href="#">Workstation Handler (WSH)</a>	A multi-contexted gateway process on a server that manages service requests from Workstation clients (that is, client processes running on remote sites).
<a href="#">Workstation Listener (WSL)</a>	A server process running on an application site that listens for and distributes connections from Workstation clients (client processes running on remote sites).
<a href="#">ULOG (User Log)</a>	A file in which error messages are stored.

## See Also

- “What Is a Single-machine Configuration” on page 3-43
- “What Are the Domains Administrative Tools” on page 3-57
- “How to Create the Configuration File for a Multiple-domain Application” on page 3-4 in *Setting Up a BEA Tuxedo Application*

# Features of a Multiple-domain Configuration

A configuration that includes more than one domain requires platform interoperability and server transparency:

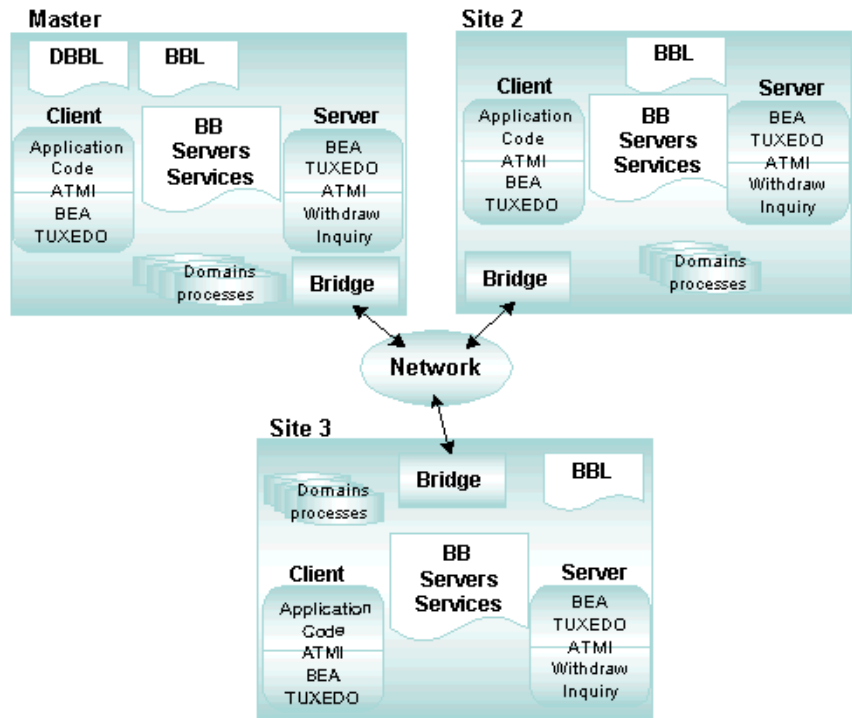
- *Platform interoperability* means that your application can rely on intermachine communications even when different machines are running different operating systems, without code customization.
- *Server transparency* means that a client can access a server without specifying its location. The locations of servers are recorded in the bulletin board and accessed as needed. As a result, servers can be moved, dropped, or added to an application dynamically, without needing to change the application itself.

## What Is a BEA Tuxedo BRIDGE

A BEA Tuxedo BRIDGE is a server, provided by the BEA Tuxedo system, for sending and receiving service requests between machines, and routing requests to local server queues.

Each BRIDGE enables a network connection to be created with every other BRIDGE in the system. Network connections are established as needed and then maintained indefinitely. BRIDGES are hidden servers, that is, they are started and stopped automatically, as needed, without an explicit configuration entry. Messages are asynchronously sent across these persistent network connections. No network connection overhead is incurred for individual messages.

Figure 3-14 Using BRIDGEs in a Multiple-machine (Distributed) Application



## See Also

- “Setting Up the Network for a Distributed Application” on page 8-1 in *Setting Up a BEA Tuxedo Application*
- “Creating the Configuration File for a Distributed Application” on page 7-1 in *Setting Up a BEA Tuxedo Application*



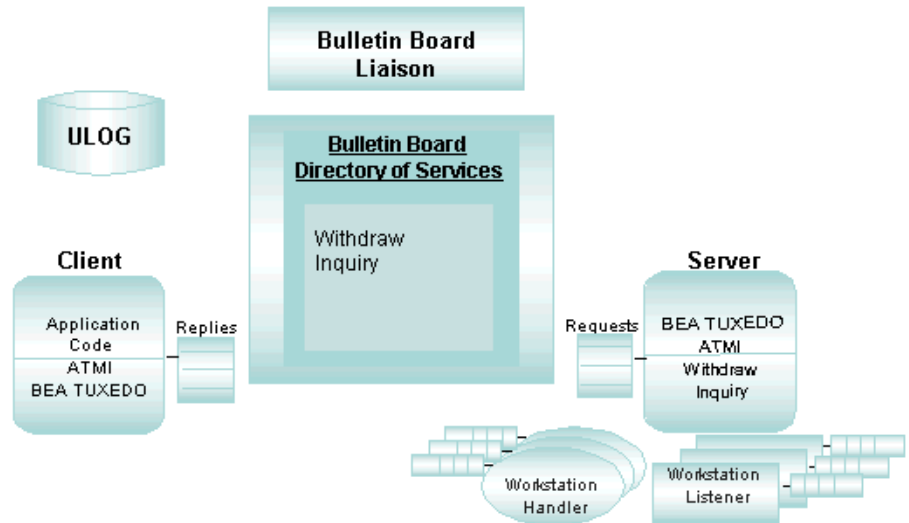
# What Is the Role of the Bulletin Board and Bulletin Board Liaison

The *bulletin board* (BB) is a memory segment in which all the application configuration and dynamic processing information is held at run time. It provides the following functionality:

- Assigns service requests to specific servers. When a service is called, the bulletin board looks up servers that offer the requested service. Based on this information, and any data-dependent routing criteria, the bulletin board places the request data on the request queue of a valid server.
- Maintains dynamic information about the *state* of an application, such as how many requests are waiting on a given server's queue and how many requests have been processed.
- Provides server location transparency, allowing an application to be developed independently of deployment. Therefore, development and deployment costs are minimized.
- Supports service name aliases, allowing multiple names to be assigned to the same service. This capability is useful for constructing interpreters, such as gateways.

The *bulletin board liaison* (BBL) is a BEA Tuxedo server that performs periodic health checks of the bulletin board and coordinates functions of all parts of the system.

Figure 3-15 Bulletin Board and Bulletin Board Liaison



## What Are Clients and Servers

- *Clients*—A client is a program that collects a request from a user and passes that request to a server capable of fulfilling it. It can reside on a PC or workstation as part of the front end of an application gathering input from users. It can also be embedded in software that reads a communication device such as an ATM machine from which data is collected and formatted before being processed by BEA Tuxedo servers.
- *Servers*—A BEA Tuxedo server is a process that oversees a set of services and dispatches services automatically for clients that request them. A service, in turn, is a function within the server program that performs a particular task needed by a business. A bank, for example, might have one service that accepts deposits and another that reports account balances. A server at this bank might receive requests from clients for both services. It is the server's job to dispatch each request to the appropriate service.

## **What Is the Distinguished Bulletin Board Liaison (DBBL)**


The *Distinguished Bulletin Board Liaison* (DBBL) is the server that makes it possible to distribute an application across multiple machines. The DBBL ensures that the bulletin board liaison (BBL) server on each machine is alive and functioning correctly. The DBBL runs on the `MASTER` machine of an application and communicates directly with all administration facilities.

The DBBL ensures that configuration and service addressing information is replicated to the bulletin board on each machine in the configuration. Servers located on remote machines are accessed through the `BRIDGE` on the local machine. Servers on the local machine are accessed directly. All local communications are performed through high performance operating system message queues. Remote communications are performed in two phases. First, service requests are forwarded to a remote machine through the (local) `BRIDGE`. Second, when a request reaches the remote machine, operating system messages are used to send the request to the appropriate server.

## **What Are the Domains Administrative Tools**

To build a multiple-domain configuration, you need to integrate your existing BEA Tuxedo application with other domains. You need to ensure interoperability across domains, preserve access to services on all domains, and accept service requests from all domains. You can perform these functions through a highly asynchronous multitasking gateway that processes outgoing and incoming service requests to or from all domains. To use the gateway, you must add entries for domain gateway groups and gateway servers to the `TUXCONFIG` file. The following illustration shows the tools provided by the BEA Tuxedo system for setting up and maintaining a multiple-domain configuration.

**Figure 3-16 Domains Administrative Tools**



Domains Tool	Description
<code>dmadmin(1)</code>	A command that allows you to configure, monitor, and tune domain gateway groups dynamically. Use this command to update the <code>BDMCONFIG</code> file while an application is running. The command acts as a front-end process that translates administrative commands to service requests to the <code>DMADMIN</code> service, a generic administrative service advertised by the <code>DMADM</code> server. The <code>DMADMIN</code> service invokes the validation, retrieval, or update functions provided by the <code>DMADM</code> server to maintain the <code>BDMCONFIG</code> file.
<code>DMCONFIG(5)</code> , <code>BDMCONFIG</code>	All Domains configuration information is stored in a <i>binary</i> file called the <code>BDMCONFIG</code> file. You can create and edit the <i>text</i> version of the Domains gateway configuration file, <code>DMCONFIG</code> , with any text editor. You can update the compiled <code>BDMCONFIG</code> file while the system is running.
<code>dmloadcf</code> and <code>dmunloadcf</code>	<code>dmloadcf</code> —reads the <code>DMCONFIG</code> file, checks the syntax, and optionally loads a binary <code>BDMCONFIG</code> configuration file <code>dmunloadcf</code> —translates the <code>BDMCONFIG</code> configuration file from binary to text format

<b>Domains Tool</b>	<b>Description</b>
DMADM(5)	A Domains administrative server that enables you to manage a Domains configuration at run time. DMADM provides a registration service for gateway groups. This service is requested by GWADM servers as part of their initialization procedure. The registration service downloads the configuration information required by the requesting gateway group. The DMADM server maintains a list of registered gateway groups, and propagates to these groups any changes made to the configuration.
GWADM(5)	A gateway administrative server that supports run-time administration of a specific gateway group. This server registers with the DMADM server to obtain the configuration information used by the corresponding gateway group. GWADM accepts requests from DMADMIN to obtain run-time statistics or to change the run-time options of the specified gateway group. Periodically, GWADM sends an “I-am-alive” message to the DMADM server. If no reply is received from DMADM, GWADM registers again. This process ensures the GWADM server always has the current copy of the Domains configuration for its group.
GWTDOMAIN(5)	A gateway process that receives and forwards messages from clients and servers in all connected domains (for TDomains).
BDMCONFIG	The binary version of the configuration file for a multiple-domain configuration.

## What Are IPC Message Queues

The BEA Tuxedo system uses IPC message queues to support communication between processes that are executed on a particular machine. IPC message queues are transient memory areas, typically provided by the underlying operating system, used for communication between clients and servers. By default, each server has its own IPC message queue on which to receive requests and replies, referred to as a Single Server, Single Queue (SSSQ). If you prefer, however, you can override the default and

assign multiple servers to read from the same queue. This arrangement is referred to as Multiple Servers, Single Queue (MSSQ). You can use both SSSQ and MSSQ sets in the same application. Servers can be assigned to either type of queue.

## **When to Use Single Server, Single Queues (SSSQ)**

To understand how SSSQ sets work, consider an analogy that can be found in your supermarket, where there may be several checkout lines. Each line is like a separate queue in which customers wait for a clerk at one register, who determines how fast that line is serviced. If a delay is introduced by one person, each subsequent person is also delayed on that line, but the delay has no effect on other lines. This scheme can be used to load balance and throttle work across several servers offering different kinds of services. Customers with relatively small requests can be processed by a server with a separate queue, thus speeding throughput by guaranteeing available cycles or registers for small requests.

## **When to Use Multiple Server, Single Queue (MSSQ) Sets**

The MSSQ scheme offers additional load balancing through IPC messaging, which is offered by the operating system. One queue is accommodated by several servers offering identical services at all times. If the server queue to which a request is sent is part of an MSSQ set, the message is dequeued to the first available server. Thus load balancing is provided at the individual queue level.

When a server is part of an MSSQ set, it must be configured with its own reply queue. When the server makes requests to other servers, the replies must be returned to the original requesting server; they must not be dequeued by other servers in the MSSQ set.

In many applications, Multiple Server, Single Queue (MSSQ) sets can play an important role. They are ideal when you need to minimize the total waiting time for services. If it is unacceptable for a service request to wait while a server capable of fulfilling that request remains idle, MSSQ sets should be used.

We recommend using an MSSQ set in the following situations:

- Service turnaround time is paramount.
- You have a reasonable number of servers (between 2 and 12).
- Servers offer identical sets of services.
- The messages involved are reasonably sized (less than 75% of the queue size).
- You can configure MSSQ sets to be dynamic so they automatically spawn and reduce servers based upon a queue load.

**Note:** For fault tolerance, you should always use MSSQ sets with two or more servers.

An MSSQ set is inappropriate when long messages are being passed to services. Long messages can cause a queue to be exhausted. When a queue is exhausted, either nonblocking sends fail or blocking sends block.

We recommend against using an MSSQ set in the following situations:

- Buffer sizes are large enough to exhaust one queue.
- You have a large number of servers. (You can compromise by using a few MSSQ sets.)
- Each server offers different services.

### Example

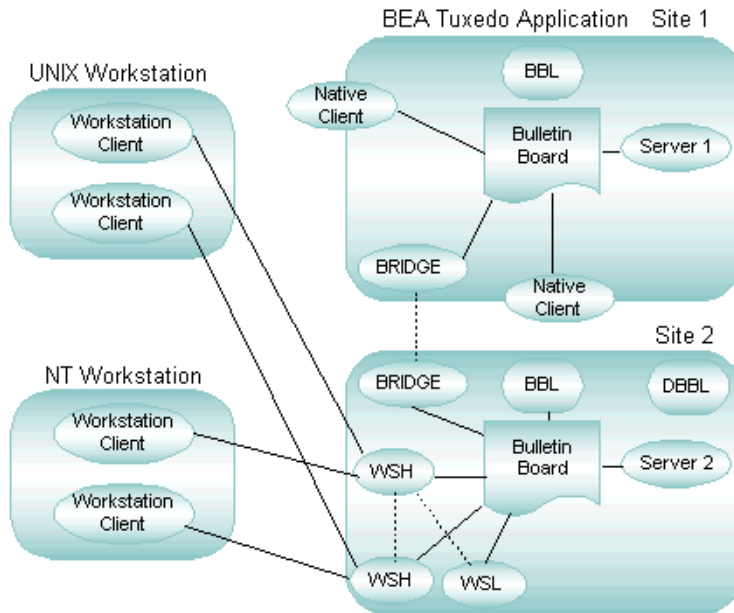
To consider how MSSQ sets work, consider an analogy that can be found in your bank, where several tellers performing identical services handle a single line of customers. The next available teller always takes the next person in line. In this scenario, each teller must be able to perform all customer services. In a BEA Tuxedo environment, all servers set up to share a single queue must offer the identical set of services at all times. The advantage of MSSQ sets is that they offer a second form of load balancing at the individual queue level.

# What Are the Workstation Handler and Workstation Listener

The Workstation component extends the availability of a native BEA Tuxedo application to clients that reside on workstations. With this component, workstations need not be within the administrative domain of the application.

The following figure shows an application with two Workstation clients (WSC). One client is running on a UNIX system workstation, while the other client is running on an NT workstation. Both WSCs are communicating with the application through the Workstation Handler (WSH) process. Initially, both joined by communicating with the Workstation Listener (WSL). The Workstation defines an environment in which clients can access the services of an application through a surrogate handler process.

**Figure 3-17 BEA Tuxedo Application with the Workstation Component**



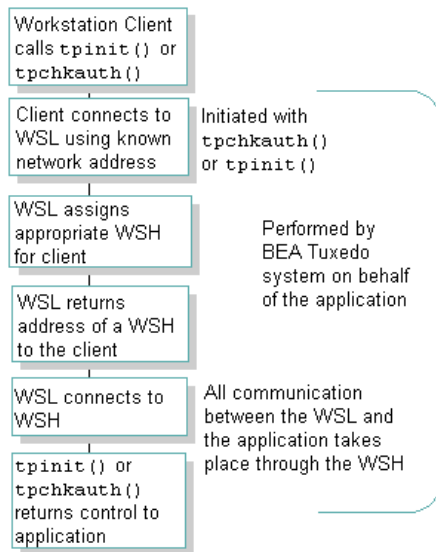


The programming environment on a Workstation is determined by the operating system of the machine. A Local Area Network (LAN) provides a connection to the administrative domain of the application, affording greater flexibility in the choice of hardware and software platforms on which you can deliver application services.

## How a Workstation Client Connects to an Application

A workstation client connects to an application in the following way.

**Figure 3-18 WSC Connecting to an Application**



## What Is the User Log (ULOG)

The user log (ULOG) is a file to which all messages generated by the BEA Tuxedo system—error messages, warning messages, information messages, and debugging messages—are written. Application clients and servers can also write to the user log. A new log is created every day and there can be a different log on each machine. However, a ULOG can be shared by multiple machines when a remote file system is being used.

The ULOG provides an administrator with a record of system events from which the causes of most BEA Tuxedo system and application failures can be determined. You can view the ULOG, a text file, with any text editor. The ULOG also contains messages generated by the `tlisten` process. The `tlisten` process provides remote service connections for other machines. Each machine, including the master machine, should have a `tlisten` process running on it.

## How Is the ULOG Created

A ULOG is created by the BEA Tuxedo system whenever one of the following activities occurs:

- A new configuration file is loaded.
- An application is booted.

## Example of a ULOG Message

The following is an example of a ULOG message.

```
121449.gumby!simpsserv.27190.1.0: LIBTUX_CAT:262: std main starting
```

A ULOG message consists of two parts: a tag and text.

The *tag* consists of the following:

- A 6-digit string (*hhmmss*) representing the time of day (in terms of hour, minute, and second)
- The name of the machine (as returned, on UNIX systems, by the `uname -n` command)
- The name and process identifier of the process that is logging the message. (This process ID can optionally include a transaction ID.) Also included is a thread ID (1) and a context ID (0).

**Note:** Placeholders are printed in the `thread_ID` and `context_ID` field of entries for single-threaded applications. (Whether an application is multithreaded is not apparent until more than one thread is used.)

The *text* consists of the following:

- The name of the message catalog
- The message number
- The BEA Tuxedo system message

---

### The Tag Indicates

- The message was written into the log at approximately 12:15 P.M.
- The machine on which the error occurred was `gumby`.
- The message was logged by the `simpsserv` process, which has a process ID of 27190.
- The thread ID is 1.
- The context ID is 0.

### The Text Indicates

- The message came from the `LIBTUX` catalog.
- The number of the message is 262.
- The message itself reads as follows: `std main starting`.

---

**Note:** For more information about a message, note its catalog name and number. With this information, you can look up the message in the appropriate catalog.

## Where the ULOG Resides

By default, the user log is called `ULOG.mmdyy` (where *mmdyy* represents the date in terms of month, day, and year) and it is created in the `$APPPDIR` directory. You can place this file in any location, however, by setting the `ULOGPFX` parameter in the `MACHINES` section of the `UBBCONFIG` file.

# 4 Integrating the BEA Tuxedo Product Family in an Enterprise System

- BEA Product Integration
- Mainframe Connectivity: Using BEA eLink
- Internet Access: Using BEA Jolt
- Developing and Managing Applications: Using BEA Manager
- Online Transaction Processing: Using the Core BEA Tuxedo System
- Enabling Scalability: Using BEA Tuxedo Domains
- Storing Messages and Service Requests: Using BEA Tuxedo /Q
- Workstation Connectivity: Using BEA Tuxedo Workstation
- Developing Client-Server Architecture Using WebLogic Enterprise
- Developing and Managing Distributed Java-based Applications: Using BEA WebLogic Server

# BEA Product Integration

The BEA product family facilitates end-to-end integration of heterogeneous hardware and software environments allowing businesses to create enterprise-wide transaction processing systems. BEA products enable companies to enjoy the benefits of robust mission-critical applications with the flexibility of distributed client/server computing. Compliant with all leading industry standards, BEA products enable developers to build, deploy, manage, and connect enterprise-wide applications on more than 70 platforms. These products also provide complete integration with market-leading application development tools, systems management solutions, and legacy applications.

Among open middleware vendors, BEA provides the only complete, end-to-end enterprise middleware solution that includes leading transaction processing components, a Web application server, and messaging products.

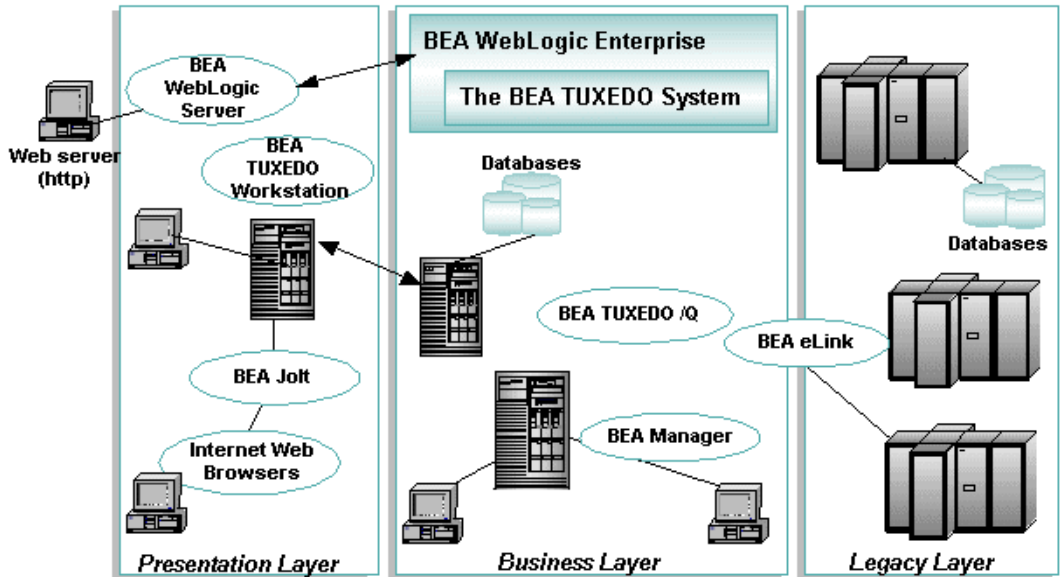
## BEA Product Suite

The BEA product suite includes the following:

- **BEA eLink**—is a family of connectivity products that seamlessly integrate BEA Tuxedo distributed applications with mainframe applications.
- **BEA Jolt**—provides Internet access and complete JAVA development support. Jolt takes requests from JAVA-enabled clients and translates them into BEA Tuxedo application calls.
- **BEA Manager**—provides tools for administration.
- *BEA Tuxedo system* comprises four components:
  - **BEA Tuxedo core product**—The most widely deployed transaction middleware for building high-performance, mission-critical, and reliable distributed applications. It provides the industry leading middleware framework for building scalable 3-tier client-server applications in heterogeneous, distributed environments.

- **BEA Tuxedo Workstation**—provides full support for clients on Windows, OS/2, Mac, UNIX, and other platforms, allowing applications to use remote clients that do not need a full BEA Tuxedo implementation.
- **BEA Tuxedo Domains**—extends the BEA Tuxedo client/server model to provide transaction interoperability across domains.
- **BEA Tuxedo /Q**—allows clients and servers to store messages or service requests for guaranteed processing.
- **BEA WebLogic Enterprise**—enables you to build, deploy, and manage component-based solutions on a proven platform. Based on the power, robustness, and demonstrated reliability of the BEA Tuxedo system, BEA WebLogic Enterprise brings together the best of industry standards, such as Common Object Request Broker Architecture (CORBA) and Enterprise Java Beans (EJB). With the power of component-based programming, these tools enable you to create a complete production-ready enterprise solution.
- **BEA WebLogic Server**—is a Java-application server that enables you to develop, integrate, deploy, and manage large-scale, distributed Web, network, and database applications.

Figure 4-1 Integration of BEA Products



## Mainframe Connectivity: Using BEA eLink

The BEA eLink family of products offers seamless transactional application-to-application interoperability between BEA Tuxedo applications and IBM CICS and IMS applications, as well as OSI transaction monitors.

This interoperability is transparent to both end users and application programmers. Native APIs and TP monitors are used on both sides of the interfaces.

### See Also

- “Components of the BEA eLink Product Suite” on page 4-5



# **Components of the BEA eLink Product Suite**

- BEA eLink for Mainframe - TCP/IP for MVS (for IMS and CICS)
- BEA eLink for Mainframe - SNA
- BEA eLink for Mainframe - OSI TP

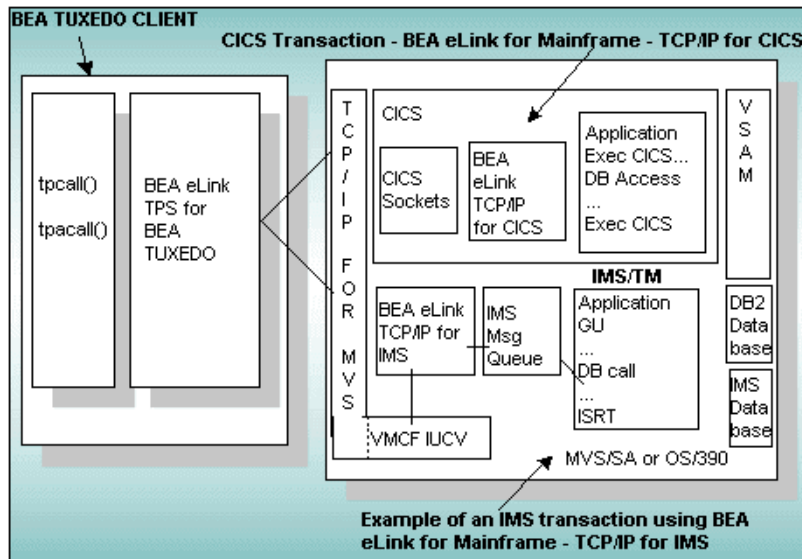
## BEA eLink for Mainframe - TCP/IP for MVS (for IMS and CICS)

The BEA eLink for Mainframe - TCP/IP gateway runs on the BEA Tuxedo system and advertises services that are available on the IBM system, either in CICS regions or under IMS.

On the mainframe, the TCP/IP for MVS communications stack is used by either the BEA eLink for Mainframe - TCP/IP for CICS or the BEA eLink for Mainframe - TCP/IP for IMS gateways to read requests from the BEA Tuxedo system and forward them to local routines.

Buffer data representation conversions and buffer layout mappings are configurable at the BEA eLink gateway. This gateway is unidirectional and nontransactional.

Figure 4-2 BEA eLink for Mainframe - TCP/IP

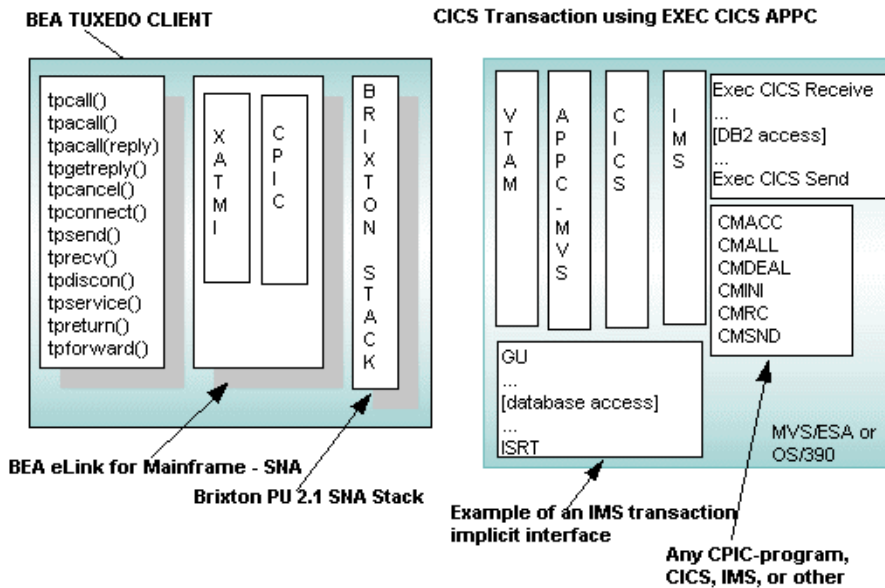


## BEA eLink for Mainframe - SNA

With BEA eLink for Mainframe - SNA, a gateway on the BEA Tuxedo system maps ATMI protocols to LU6.2/PU.1 protocols. This mapping allows routines on the mainframe to use CPIC or APPC functions to receive data from the BEA Tuxedo system and forward that data to local routines, or to initiate requests to services in the BEA Tuxedo system.

Buffer data representation conversions and buffer layout mappings are configurable at the BEA eLink gateway. BEA eLink for Mainframe - SNA is fully transactional and bidirectional.

**Figure 4-3 BEA eLink for Mainframe - SNA**

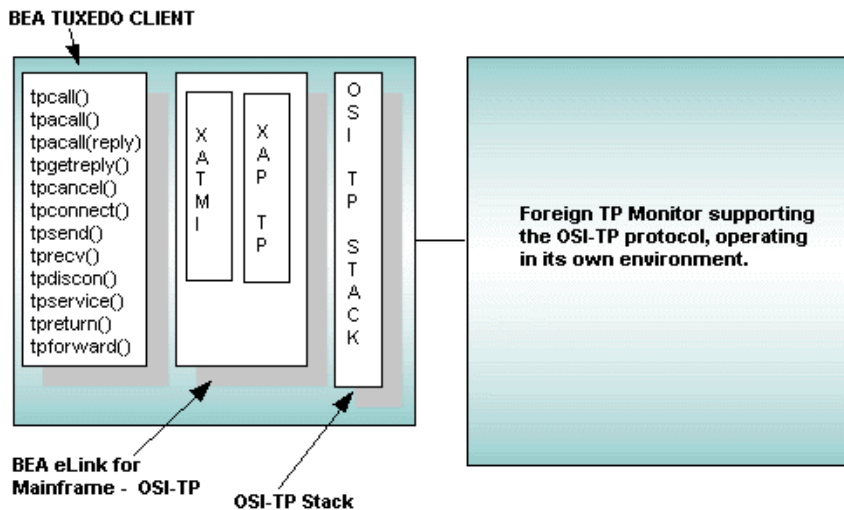


## BEA eLink for Mainframe - OSI TP

The BEA eLink for Mainframe - OSI TP product is similar to the BEA eLink for Mainframe - SNA product.

Both are built on the BEA Tuxedo Domains architecture, and provide transparent transactional bidirectional access through gateway processes, which translate between the XATMI protocol and the protocol of the foreign TP system which, in this case, is OSI TP.

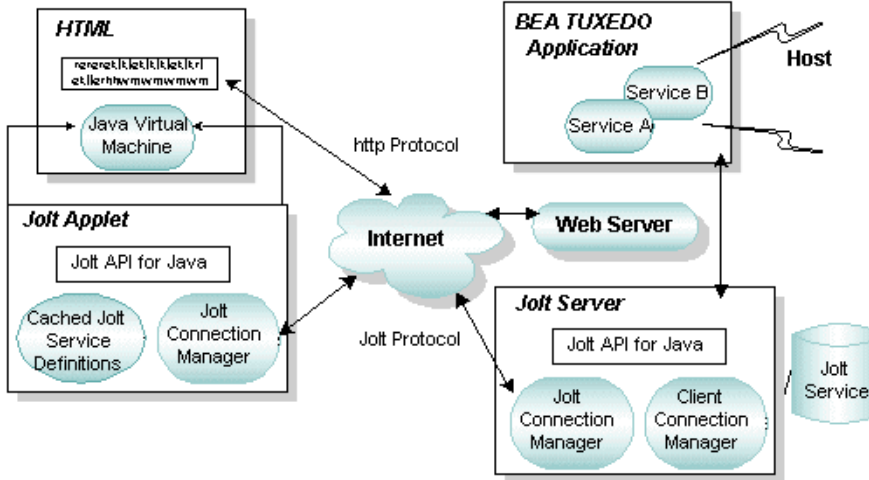
Figure 4-4 BEA eLink for Mainframe - OSI TP



# Internet Access: Using BEA Jolt

BEA Jolt allows BEA Tuxedo clients to be written in the Java language, thus making such clients available through the Internet or intranets. The product consists of two main components: the Jolt Class Library and the Jolt Repository. Both components can be used with BEA Tuxedo applications that are not written in Java. With BEA Jolt, you can create secure, scalable transactions, over the Internet, between clients and servers.

Figure 4-5 BEA Jolt



## See Also

- “Components of BEA Jolt” on page 4-10

# Components of BEA Jolt

BEA Jolt consists of a Java Class Library and an API between Java clients and the BEA Tuxedo system. It provides several components for creating Java-based client programs that access BEA Tuxedo services and for enabling secure, reliable access to servers inside corporate firewalls:

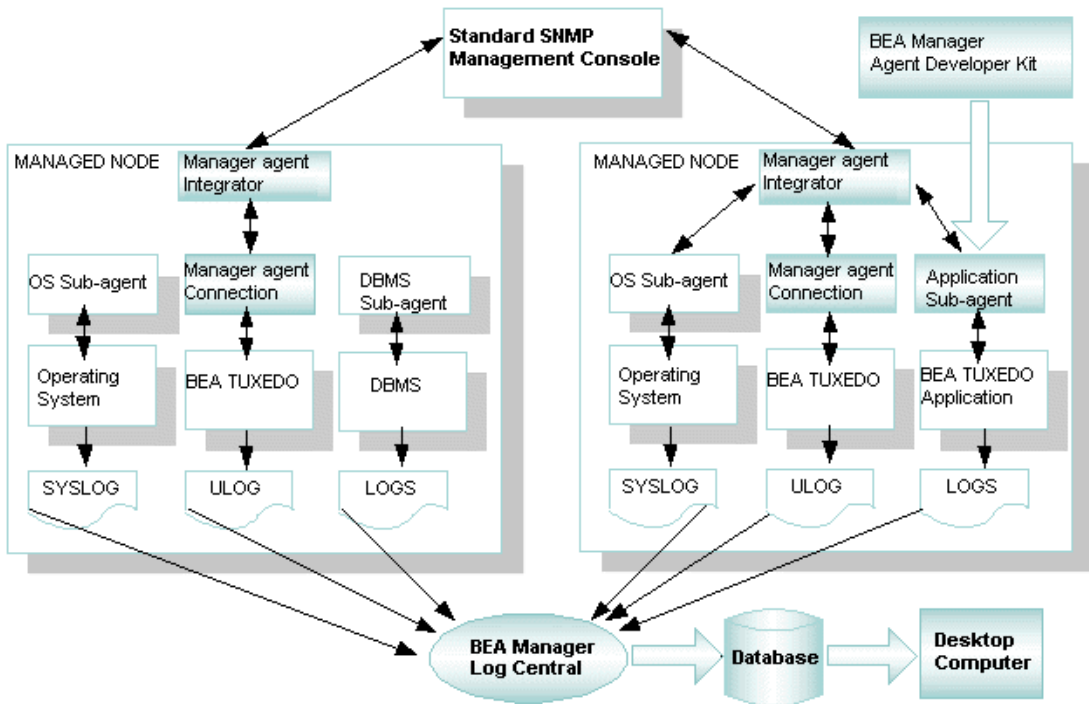
- *Jolt Server*—One or more Jolt servers listen for network connections from clients, translate Jolt messages, multiplex multiple clients into a single process, and submit and retrieve requests to and from BEA Tuxedo applications running on one or more BEA Tuxedo servers.
- *Jolt Class Library for Java*—The Jolt Class Library consists of Java class files that implement the Jolt API. These classes enable Java applications and applets to invoke BEA Tuxedo services. The Jolt Class Library provides functions for managing, retrieving, and invoking communication attributes, notifications, network connections, transactions, and services.
- *Jolt Repository*—A central Jolt Repository contains definitions of BEA Tuxedo system services. These Repository definitions are used by Jolt at runtime to access BEA Tuxedo services. Using the Repository Editor, you can test new and existing BEA Tuxedo services independently of the client applications. You can export services to a Jolt client application or unexport services by hiding the definitions from the Jolt client.
- *Jolt Internet Relay*—The Jolt Internet Relay is a component that routes messages from a Jolt client to a Jolt Server Listener (JSL) or Jolt Server Handler (JSH). This tool eliminates the need for the JSH and the BEA Tuxedo system to run on the same machine as the Web server. The Jolt Internet Relay consists of the Jolt Relay (JRLY) and the Jolt Relay Adapter (JRAD).

The separation of BEA Jolt into these components permits the transactional and Internet components of client/server applications to be implemented separately with the security and scalability required for large-scale Internet and intranet services.

# Developing and Managing Applications: Using BEA Manager

BEA Manager is an integrated set of software products that provides a complete environment for developing, managing, integrating, and deploying BEA Tuxedo applications using industry standard SNMP technology.

Figure 4-6 BEA Manager



# BEA Manager Components

**Table 4-1 Purpose of BEA Manager Components**

<b>BEA Manager Component</b>	<b>Purpose</b>
BEA Manager Agent Connection	Exposes the BEA Tuxedo Management Information Base (TMIB) to a standard SNMP Management Console (for example, OpenView, NetView, Tivoli, BMC). Information can then be viewed and managed easily through the console as a consolidated picture of the enterprise environment.
BEA Manager Agent Integrator	Solves the problem caused by running multiple SNMP agents on the same managed machine. It allows multiple SNMP agents and subagents from any vendor to operate on the same machine and to appear as a single SNMP agent to the SNMP Management Console.
BEA Manager Agent Development Kit	Enables developers to build their own SNMP agents to monitor and control BEA Tuxedo (and non-Tuxedo) applications, and any other specialized software or hardware.
BEA Manager Log Central	A centralized log management system that consolidates log messages from distributed machines to a single machine for viewing, filtering, and managing. You can select the events to be sent to an SMNP Management Console.

## Components of the BEA Tuxedo Product

- [BEA Tuxedo core product](#)
- [BEA Tuxedo Domains](#)
- [BEA Tuxedo /Q](#)
- [BEA Tuxedo Workstation](#)

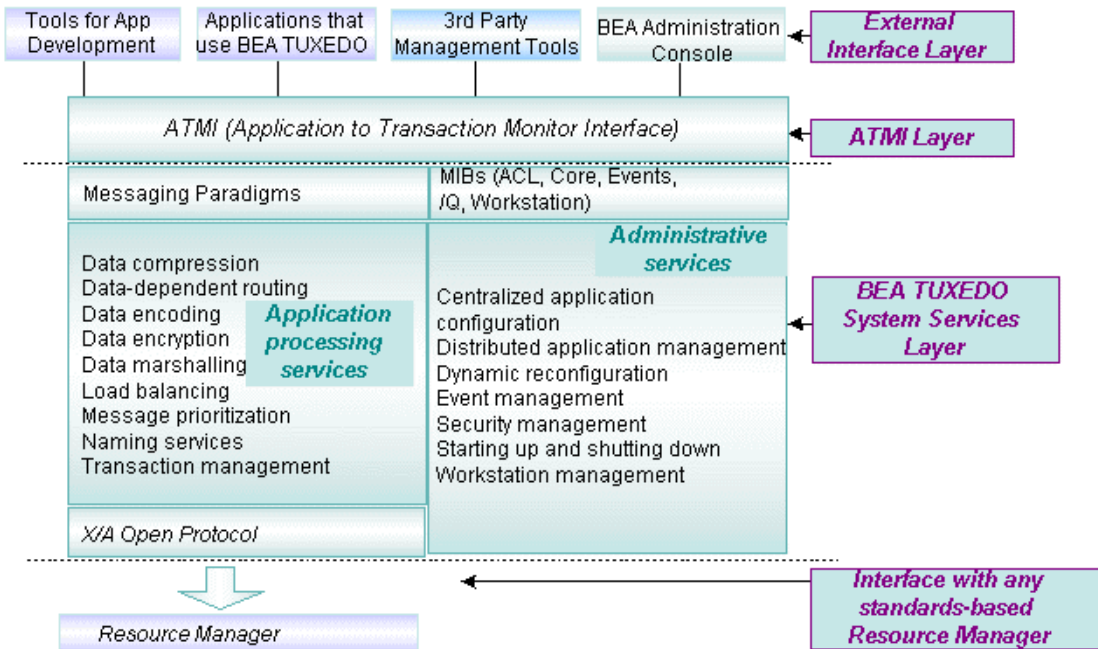


# Online Transaction Processing: Using the Core BEA Tuxedo System

The BEA Tuxedo system is a development platform that enables you to create applications that mix and match hardware platforms, databases, and operating systems to fit your business needs. The BEA Tuxedo system provides a foundation for client/server architecture, request/response and conversational communications interfaces, transaction support, and administration for a distributed application.

The BEA Tuxedo system provides all the features and benefits of a high-end Online Transaction Processing (OLTP) system, including scalability, high performance, mission-critical reliability, and standards support.

**Figure 4-7 Architecture of a Basic BEA Tuxedo System**



As shown in this illustration, the BEA Tuxedo system contains the following parts.

Architectural Part	Description
External interface layer	This layer consists of interfaces between the user and the system. It includes both tools for application development, such as Simple Network Management Protocol (SNMP) agents, and tools for administration, such as the BEA Administration Console. The BEA Administration Console and SNMP agents can interact with standard management consoles. Thus a user can manage a BEA Tuxedo system and a network configuration from one console. In addition, application architects and developers can build their own administrative tools or application- or market-specific tools on top of the MIB.
<a href="#">MIB (Management Information Base)</a>	The Management Information Base (MIB) is an interface that enables users to program and administer a BEA Tuxedo system easily. MIB operations enable you to perform all management tasks (monitor, configure, tune, and so on). The MIB allows you to perform one task to one object at a time or to build tool kits with which you can batch tasks and/or objects. (For information regarding the different parts of the MIB, see “Available BEA Tuxedo System MIBs” on page 3-3.)
<a href="#">Administration Console</a>	A Web-based graphical user interface for managing BEA Tuxedo applications. The interface allows you to enter and modify data in the MIB. The BEA Administration Console makes these tools available through a Web browser. The server-side components of the BEA Administration Console reside on one of the machines in your BEA Tuxedo domain. To use the Console, you must enter the URL of the server and download a set of Java applets, which implement the Console. The Console enables any user with a supported browser to administer the BEA Tuxedo system.
<a href="#">ATMI (Application to Transaction Monitor Interface)</a>	The interface between an application program and the BEA Tuxedo system. The ATMI and the BEA Tuxedo system implement the X/Open DTP model of transaction processing. An abstract environment, the ATMI supports location transparency and hides the details of implementation. As a result, programmers are free to configure and deploy BEA Tuxedo applications without modifying the application code.

Architectural Part	Description
BEA Tuxedo Services ( <a href="#">administrative services</a> and <a href="#">application processing services</a> )	Services and/or capabilities common to the BEA Tuxedo system infrastructure for developing and administering applications. The application processing services available to developers include: transactions, messaging paradigms, type validation, load balancing, data-dependent routing, service prioritization, data encoding, marshalling, and compression, and reliable queueing. The administrative services include: distributed transaction processing, security management, service naming, distributed application administration, centralized application configuration, dynamic reconfiguration, and domains partitioning.
Resource Manager	A software product in which data is stored and available for retrieval through application-based queries. The resource manager (RM) interacts with the BEA Tuxedo system and implements the XA standard interfaces. The most common example of a resource manager is a database. Resource managers provide transaction capabilities and permanence of actions; they are the entities accessed and controlled within a global transaction.

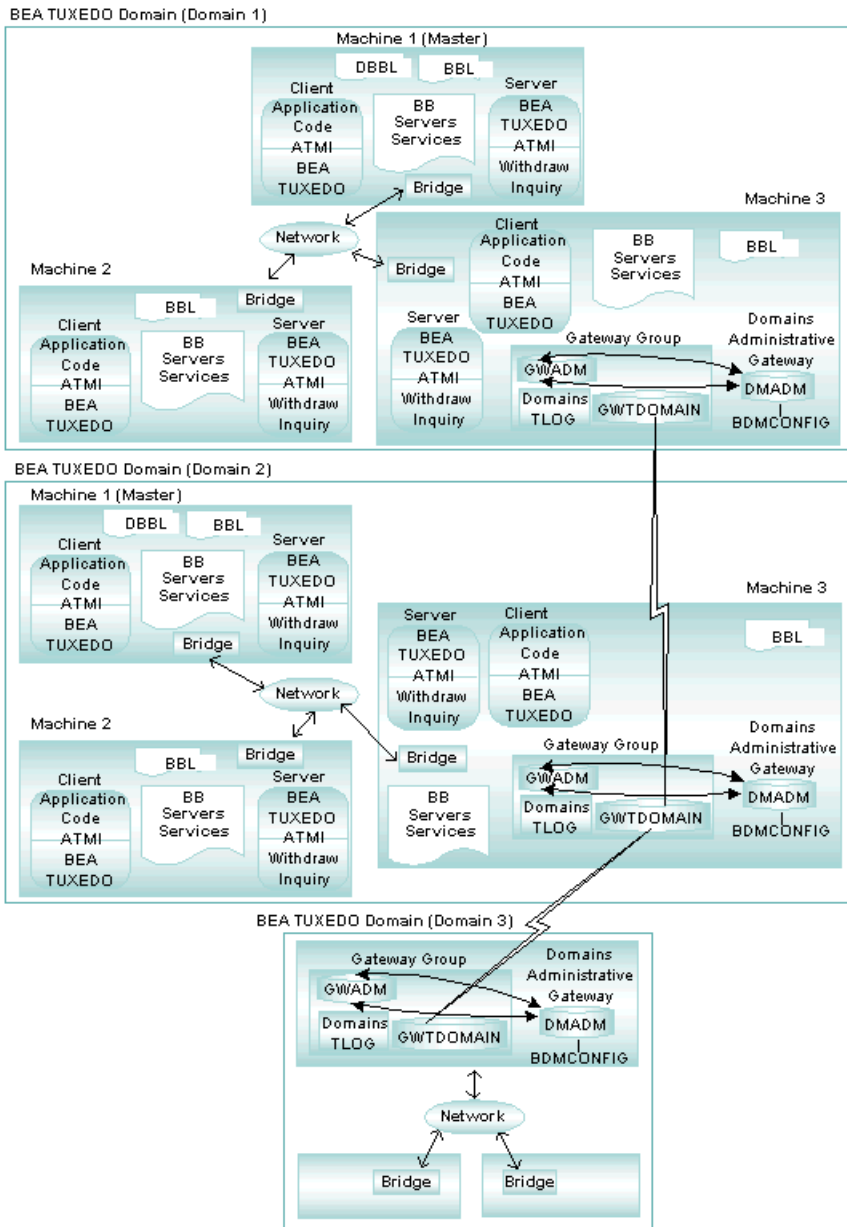
## See Also

- “BEA Tuxedo System Architecture” on page 2-1

# Enabling Scalability: Using BEA Tuxedo Domains

Domains extends the BEA Tuxedo system client/server model to provide transaction interoperability across TP domains. This extension preserves the client/server model and the ATMI interface by making access to services on a remote domain (and service requests from a remote domain) transparent to both the application programmer and the user. Domains makes this possible through a highly asynchronous multitasking gateway that processes service requests to and from remote domains.

Figure 4-8 Multiple-domains Environment



## Features of BEA Tuxedo Domains

BEA Tuxedo Domains provides the following features:

- *Aliasing capability*—An administrator can map the service names used by a remote application to the service names used by a local application, allowing easy integration of applications that use different naming schemes.
- *Availability*—Multiple network addresses can be defined for a remote domain. An administrator can specify a backup domain for a set of services.
- *Scalability and modular growth*—Application programmers can structure their applications for modularity, isolation of failures, and independent growth. Interoperation with other TP applications can be achieved easily by adding a description of the interfaces (that is, services) between a remote application and the Domains configuration.
- *Security*—The access control list (ACL) facility restricts access to local services from a particular set of remote domains. The security feature also provides a structuring capability for defining different views of the exported services available to remote domains.
- *Transparency and independence*—Applications are completely unaware of service distribution. Client application programmers need not know the implementation changes made to a service, the location of a service, or network addresses. A service can be available on the same machine as a client, on another machine in the local domain, or on a remote domain.

## See Also

- “What Is a Domain” on page 4-18
- “What Is a Domains Gateway” on page 4-18
- “What Are Domain Gateway Types” on page 4-19
- “BEA Tuxedo Domains Components” on page 4-20

# What Is a Domain

A domain consists of a BEA Tuxedo application running one or more business applications. A single domain is defined in one configuration file and is administered as a single entity.

# What Is a Domains Gateway

A Domains gateway (`GWTDOMAIN`) is a server provided by the BEA Tuxedo system that enables access to and from remote domains. A Domains gateway runs as a Multiple-Server, Single-Queue set (`MSSQ`) in which each gateway uses a common request queue and has its own reply queue. In addition, Domains provides two gateway administrative servers. `GWADM` enables run-time administration of the Domains gateway group. `DMADM`, a Domains administrative server, enables run-time administration of Domains configuration information.

Domains gateways support the following functionality:

- *Administration*—A gateway can be booted or shut down exactly as any other BEA Tuxedo server. Run-time administration is provided through an administrative server, `DMADM`. Using `DMADM`, application administrators can change the domain configuration file and tune the performance of a gateway group.
- *ATMI conversational model*—Application programs can establish conversations with programs running in another domain. Remote domains can establish conversations with conversational services offered by local servers.
- *ATMI request/response model*—Application programs using the BEA Tuxedo system can request services from applications running in another domain. Also, remote applications can request services from local servers. No changes are required to the application programs.
- *Multi-domain interaction*—Gateways can communicate with multiple domains of the same type.

- *Multi-network support*—Gateways can communicate with other domains through a variety of networks such as Ethernet and Novell. However, a gateway is limited by the capabilities of the networking library to which it is linked. In other words, a gateway typically supports a single type of network.
- *Transaction management*—Application programs can interoperate with other domains within a transaction. The gateway coordinates the commitment or rollback of transactions running across domains.
- *Typed buffer support*—Gateways can perform encoding and decoding operations for all typed buffers defined by the application.

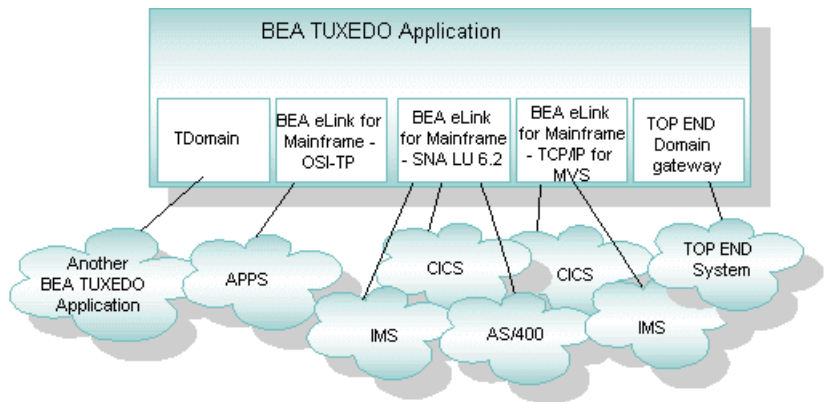
## What Are Domain Gateway Types

Communication among domains is managed through a set of processes called gateways. The BEA Tuxedo system offers the following types of domain gateways to accommodate different communication protocols:

- *BEA Tuxedo (TDomains)*—Domains provides interoperability between two or more BEA Tuxedo applications through a specially designed TP protocol that flows over network transport protocols such as TCP/IP.
- *BEA eLink for Mainframe-OSI TP*—provides interoperability between BEA Tuxedo applications and other transaction-processing applications that use the OSI TP standard. OSI TP is a protocol for distributed transaction processing defined by the International Standards Organization (ISO).
- *BEA eLink for Mainframe-SNA*—allows BEA Tuxedo clients and servers to operate with clients and servers in an MVS/CICS or MVS/IMS environment in remote SNA domains. It also connects to multiple SNA networks from the local BEA Tuxedo domain.
- *BEA eLink for Mainframe-TCP/IP for MVS (for CICS)*—provides gateway connectivity that makes it possible for non-transactional tasks within BEA Tuxedo regions to access services provided by CICS application programs and vice-versa. It enables a BEA Tuxedo domain to communicate, via the TCP/IP network transport protocol, to a CICS environment. (A BEA Tuxedo domain is the BEA Tuxedo application defined in a single configuration file.)

- *BEA eLink for Mainframe-TCP/IP for MVS (for IMS)*—provides gateway connectivity enabling transparent communications between client and server transactions in an IMS system and a BEA Tuxedo domain, a CICS system, or another IMS system.
- *BEA TOP END Domain Gateway (TEDG)*—provides interoperability between TOP END systems and BEA Tuxedo domains. The TEDG supports fully transactional bidirectional message passing and queuing between TOP END and BEA Tuxedo systems.

**Figure 4-9 Domain Gateway Types**



## BEA Tuxedo Domains Components

BEA Tuxedo Domains includes the following components:

- *A domain gateway*—A BEA Tuxedo system program that enables access to and from remote domains. There are four gateways:
  - *TDomains* provides interoperability between two or more BEA Tuxedo applications through a specially designed TP protocol that flows over network transport protocols such as TCP/IP.



- *BEA eLink for Mainframe-OSI TP* provides interoperability between BEA Tuxedo applications and other transaction processing applications that use the OSI TP standard.
- *BEA eLink for Mainframe-SNA* enables BEA Tuxedo clients and servers to operate with clients and servers in an MVS/CICS, MVS/IMS, or AS/400 environment in remote SNA domains. It also connects to multiple SNA networks from the local BEA Tuxedo SNADOM.
- *BEA TOP END Domain Gateway (TEDG)* provides interoperability between TOP END systems and BEA Tuxedo domains.
- *A gateway administrative server (GWADM(5))*—A BEA Tuxedo program that enables run-time administration of a particular domain gateway group.
- *A domain administrative server (DMADM(5))*—A BEA Tuxedo program that enables run-time administration of the configuration information required by domain gateway groups.
- *An administrative interface*—An interface for the configuration and run-time administration of the information required by domain gateways for interoperation with other domains.


## Storing Messages and Service Requests: Using BEA Tuxedo /Q

The BEA Tuxedo /Q extension allows clients and servers to store messages or service requests for guaranteed processing. The system guarantees requests are sent through the transaction protocol to ensure safe storage. The administrative functions of BEA Tuxedo /Q provide the administrator with a great deal of flexibility in establishing and managing queues, and configuring system servers provided with the queued message facility.

## Using the Message Queuing Server

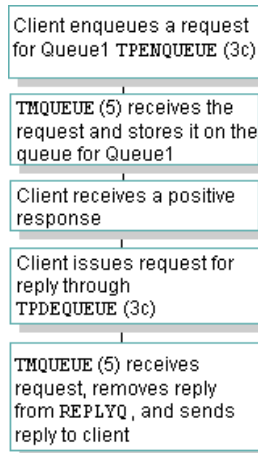
A *message queuing* server (TMQUEUE) allows for transparent enqueueing and dequeuing of messages. The following illustration shows how TMQUEUE works.

**Figure 4-10 Queueing Messages Using TMQUEUE**



qmadmin(1) is used to create the QUEUE SPACE (a set of queues managed by a queue manager) with *one queue* (Queue1). The following flowchart explains how messages are queued and dequeued.

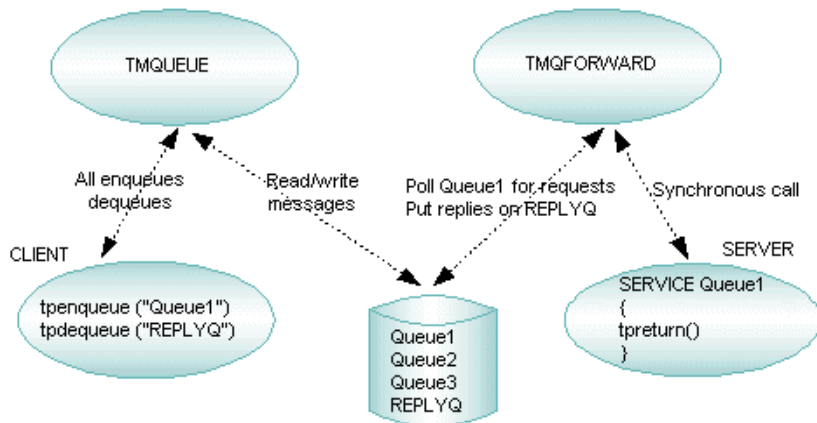
**Figure 4-11 Process of Enqueueing and Dequeueing Messages**



## Storing and Forwarding Messages

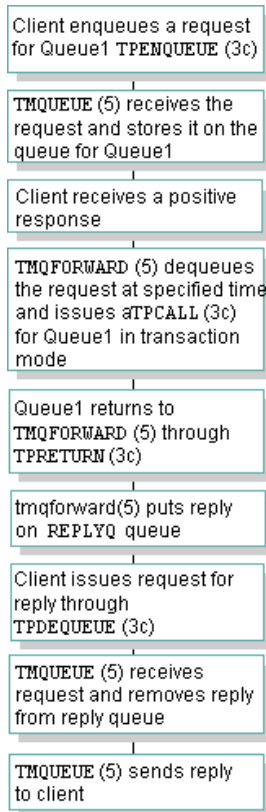
A *forwarding* server (TMQFORWARD) dequeues messages and forwards them to the appropriate servers for processing. This allows transparent processing of enqueued messages by BEA Tuxedo system servers unaware whether the incoming message was sent as a normal request/response message or through the stable queue. A response by the server is automatically enqueued to an associated reply queue for each message. The following illustration shows how the TMQFORWARD server works.

**Figure 4-12 Storing and Forwarding Messages using TMQFORWARD**



qmadmin(1) is used to create the QUEUE SPACE with *four queues*: Queue 1, Queue2, Queue3, and REPLYQ. The following flowchart explains how messages are stored and forwarded.

**Figure 4-13 Process of Storing and Forwarding Messages**



The TMQFORWARD(5) server is needed only if queued messages require a service call. For example, a queue may be used (on a BEA Tuxedo client or server) for interprocess communication in which one process places the message on the queue and another removes it. (The call to tpdequeue(3c) and the use of a reply queue are also optional.)


## BEA Tuxedo /Q Capabilities

- /Q makes it possible for an application, within a global transaction, to enqueue a message in a stable storage area for processing at a later time.
- /Q provides administrators of BEA Tuxedo applications with functions for processing messages on queues. The ATMI interface includes one function, `tpenqueue(3c)`, that allows clients and servers to store messages on a particular queue, and another, `tpdequeue(3c)`, that allows clients and servers to retrieve messages from a particular queue. The `qmadmin(1)` command provides administrative control of queues and queued messages.
- Requests can be dequeued on the basis of LIFO, FIFO, or time. Within LIFO and FIFO, requests can be dequeued by priority.
- The queue space is a resource manager that complies with X/Open's XA interface to ensure data integrity in a distributed environment.
- When the BEA Tuxedo system is used in an environment in which a machine, server, or resource is sometimes unavailable or unreliable, as in a wide-area network, /Q can provide continuous processing, storing messages until the resources necessary to process those messages are available.
- In batch processing of a potentially long-running transaction, you are guaranteed that a message will be processed eventually so you do not have to wait until the transaction is completed.
- /Q can be used for work flow provisioning such that each step generates a queued request to perform the next step in a process. Mechanisms such as data-dependent routing and priority handling are kept intact for queue-based requests and replies.
- /Q can be combined with BEA Tuxedo Workstation to enqueue and dequeue messages from Workstation clients. The interface for this combination is available in both the C and COBOL programming languages.

# Workstation Connectivity: Using BEA Tuxedo Workstation

The Workstation component of the BEA Tuxedo system allows application clients to be located on remote sites that are not part of a BEA Tuxedo domain. These remote sites do not support an administration server, an application server, or a bulletin board. All communication between such a client and an application takes place over a network. The following illustration shows a Workstation environment.

**Figure 4-14 Connecting Clients Using Workstation**



Workstation extends the benefits of the BEA Tuxedo system to the desktop at run-time, and for programming and administering BEA Tuxedo applications. Run-time advantages include: more dynamic connectivity, less administrative overhead, greater security as it keeps clients off the server systems, greater server utilization as it offloads CPU cycles and decreases process context switches, and a smaller footprint.

## See Also

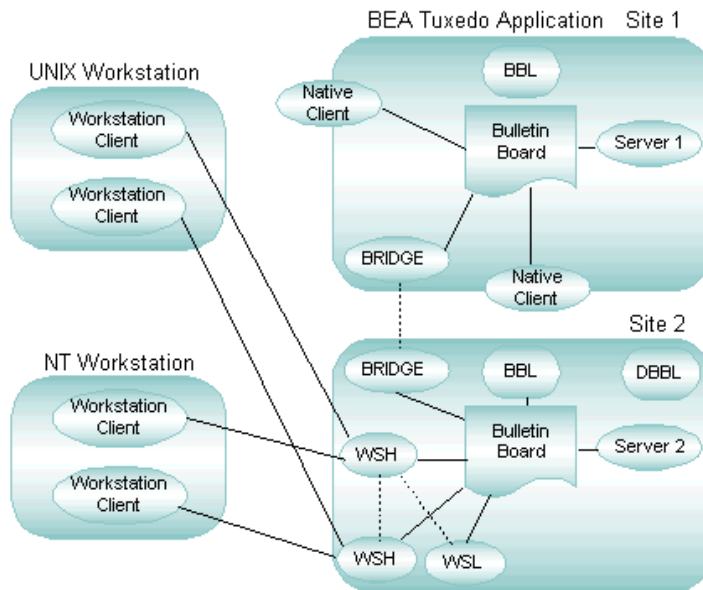
- “Workstation Components” on page 4-27

# Workstation Components

The BEA Tuxedo Workstation components are:

- *Workstation Clients*—Clients use BEA Tuxedo Workstation to connect to the BEA Tuxedo system. The BEA Tuxedo system supports multiple Workstation clients running simultaneously on Windows, Windows NT, UNIX, and VMS platforms. The ATMI is available to clients supported by Workstations.
- *Workstation Handler (WSH)*—A WSH manages one or more connections between Workstation clients and native BEA Tuxedo servers, makes surrogate service requests, manages transactions, and returns replies. WSH processes are started and stopped by the WSL.
- *Workstation Listener (WSL)*—A WSL accepts connection requests from Workstation clients and assigns connections to a Workstation Handler. It also manages the pool of Workstation Handler processes, starting and stopping them in response to load demands.

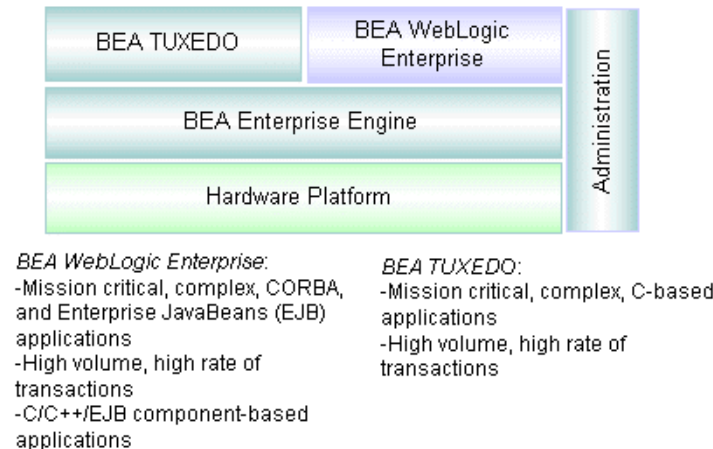
**Figure 4-15 Workstation Components**



# Developing Client-Server Architecture Using WebLogic Enterprise

WebLogic Enterprise (WLE) provides businesses and organizations that depend on mission-critical applications with the advantages of CORBA (Common Object Request Broker Architecture)-compliant and Enterprise Java Beans (EJB) programming models combined with the power, robustness, and proven reliability of the BEA Tuxedo system. The WLE deployment infrastructure, based on the BEA Tuxedo system, delivers secure, transactional, and distributed applications in a managed environment. Traditional procedural OLTP functions are provided through the WLE ATMI, which is based on the BEA Tuxedo ATMI.

**Figure 4-16 WebLogic Enterprise**



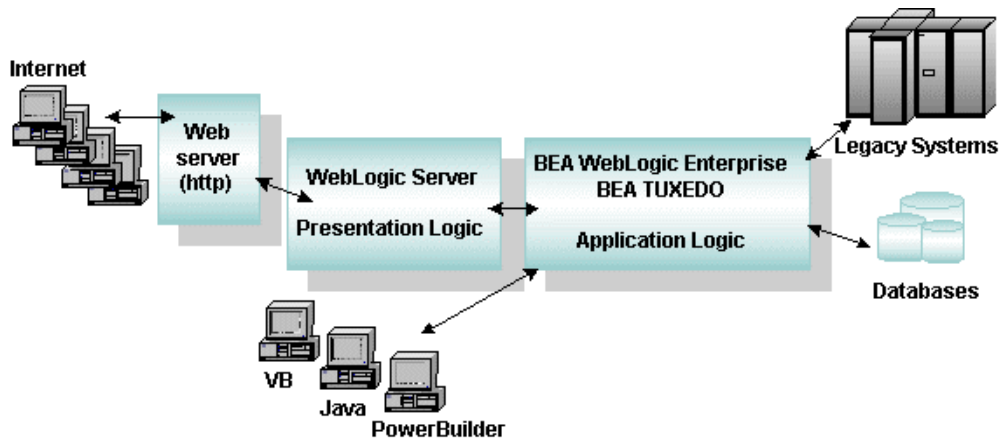
BEA WebLogic Enterprise provides companies with a reliable way to develop e-commerce applications that can be distributed across thousands of local or geographically dispersed computers. BEA WebLogic Enterprise delivers advanced scalability, reliability, and expandability in a standards-based product. In a single package, BEA WebLogic Enterprise offers customers wide flexibility in development models and programming languages, including CORBA, Java, C++, C, COBOL, Java 2 Enterprise Edition (J2EE), and Enterprise Java Beans (EJBs). WLE also enables companies to deploy their applications via a server that can scale to reliably support



millions of online e-commerce customers. Customers can mix and match CORBA objects, BEA Tuxedo services, and EJB components in the same application. In addition, BEA WebLogic Enterprise provides a platform for end-to-end Web-to-mainframe integration, helping businesses to offer a more competitive set of products and services by harnessing their IT infrastructures in e-business initiatives, instead of standalone Web sites.

Because BEA WebLogic Enterprise is built on a proven, reliable BEA Tuxedo engine, this product provides all the features necessary to enable companies to deliver secure e-commerce solutions rapidly.

**Figure 4-17 WebLogic Server and WebLogic Enterprise**

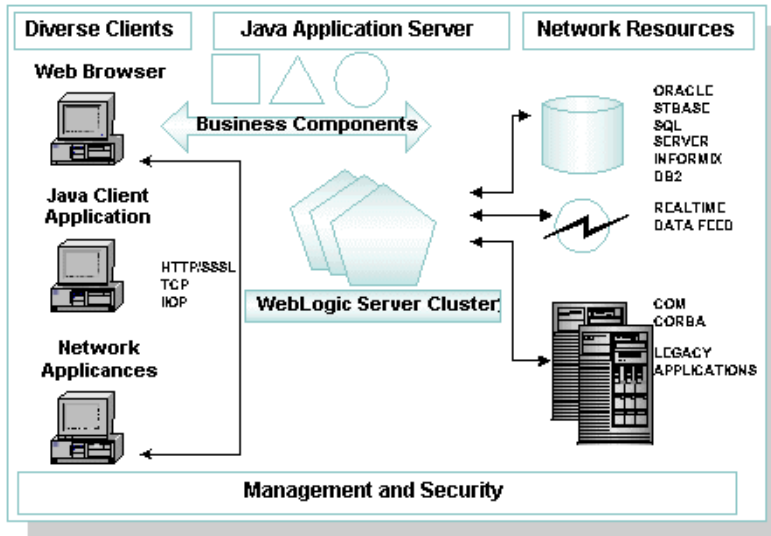


# Developing and Managing Distributed Java-based Applications: Using BEA WebLogic Server

BEA WebLogic Server is a Java application server for developing, integrating, deploying, and managing large-scale, distributed Web, network, and database applications. BEA WebLogic Server makes it possible to build portable, scalable applications that interoperate seamlessly with other applications and systems.

The principal advantage of BEA WebLogic Server is the power it brings to the software development cycle. Low-level network services—for example, sockets—are isolated and handled automatically by higher level BEA WebLogic Server facilities—such as its RichSockets™ and connection pool design. High-level components such as WebLogic Beans and WebLogic/JDBC provide enhanced JavaBeans and JDBC development services for the WebLogic environment.

Figure 4-18 BEA WebLogic Server



## **WebLogic Server Implementations**

BEA WebLogic Server offers the following implementations:

- The BEA WebLogic Server is the heart of a distributed application suite; it can manage a broad range of software and hardware services within a distributed computing framework.
- BEA WebLogic Server/JDBC is a Type 3 implementation of JDBC for use with Java applets or applications. BEA WebLogic Server/JDBC provides a subset of the Enterprise Java Standard services available in the BEA WebLogic Application Server. These services include security, HTTP servlet support for server-side programming, name services, and access control lists. BEA WebLogic Server/JDBC comes with one of the industry-leading WebLogic two-tier JDBC drivers for Oracle, Sybase, or MS SQL Server.
- BEA WebLogic Server JDBC Drivers are the most widely used JDBC drivers in the industry. They include: Type 2 (2-tier) JDBC implementations for Oracle, Sybase, and Microsoft SQL Server databases, and Type 4 implementations for Microsoft SQL Server and Informix.

## **Advantages of Using BEA WebLogic Server**

- Comprehensive support for the Enterprise Java Standards protects your investment and makes it possible to build portable, scalable applications that interoperate seamlessly with other applications and systems.
  - Most Enterprise JAVA APIs (10 of 12) are fully implemented.
  - The most comprehensive implementation of the Enterprise JavaBeans 1.0 specification, including session and entity beans, is provided.
- The critical front-end Web component of the BEA end-to-end enterprise middleware solution is provided.

## **4** *Integrating the BEA Tuxedo Product Family in an Enterprise System*

---