



TIBCO Administrator™

User's Guide

*Version 5.12.0
September 2021*



Contents

Figures	9
Tables	13
Chapter 1 Introduction	15
Overview	16
User Management	16
Resource Management	17
Application Management	18
TIBCO Administration Domain	19
Administration Server	20
Tomcat Web Server	21
TIBCO Hawk Agent	22
TIBCO Runtime Agent	22
Utilities, Plug-ins and Modules	23
TIBCO Domain Utility	23
Scripting Deployment Utility	23
Command Line Utilities	24
TIBCO Enterprise Message Service Plug-in	26
TIBCO BusinessWorks Manual Work Module	28
UDDI Servers Module	28
Chapter 2 Starting TIBCO Administrator	29
Starting the TIBCO Administrator GUI	29
Stopping the Administration Server	30
Starting TIBCO Administrator on Microsoft Windows	32
Starting the Administration Server and TIBCO Hawk Agent	32
Starting the TIBCO Administrator GUI	32
Starting TIBCO Administrator on UNIX	34
Starting the Administration Server and TIBCO Hawk Agent	34
Starting the TIBCO Administrator GUI	34
Stopping the Administration Server	35
Chapter 3 Integrating TIBCO Administrator with an LDAP Directory Server	36
Overview	37

Supported LDAP Directory Servers	38
Features	38
Limitations	41
Managing LDAP Users and Group-synchronized Roles	45
Automatically Creating a Role for Each LDAP Group	47
Choosing Specific LDAP Groups to Synchronize	47
Adding a Local User to an LDAP Integrated Domain	48
Adding a Local Role to an LDAP Integrated Domain	49
Filtering LDAP Users and Groups to Integrate	51
Selecting LDAP Groups to Synchronize in TIBCO Administrator	51
Setting the Maximum LDAP Objects to Return After a Search	54
Pre Loading User Objects	55
LDAP Synchronization Optimistic Option	56
Prerequisites for Using the Optimistic Option	56
Chapter 4 Managing Users and Roles	57
User Management Overview	58
Using the Guest Role	58
Searching for Users	58
Selecting Items	59
Managing Access Rights	60
Adding Users	64
Assigning Role Membership to Users	65
Add Roles Dialog	65
Removing Role Membership for a User	67
Edit Roles Dialog	67
Assigning Permissions to Users	68
Changing or Resetting Passwords	69
Deleting Users	71
Renaming Users	72
Changing Domain Administrator User Credentials	73
Managing the Password Policy for an Administration Domain	75
Modifying the Password Policy	75
Removing the Password Policy	76
Password Policy Elements	76
Creating a Role	81
Adding or Removing a User from a Role	82
Removing a Child Role from a Parent Role	83
Assigning Permissions to Roles	84

Users Dialog	86
New User Dialog	87
General Tab	87
Role Membership Tab	87
Permissions Tab	88
Roles Dialog	89
New Role Dialog	90
General Tab	90
Members Tab	90
Roles Tree Tab	91
Permissions Tab	92
Security Tab	92
Profile Dialog	93
Chapter 5 Granting Security Access to Objects	94
Security Overview	95
Security Console Tree	95
Granting Super User Access	98
Granting Access to an Object	99
Managing Concurrent Access	101
Edit Security Dialog	102
Security Dialog	103
Chapter 6 Managing Installed Software and Machines	106
Resource Management Overview	107
Working With Application Domains	108
Customizing the Installed Software Display	111
Customizing the Machines Display	112
Disabling and Enabling Installed Software	113
Adding Custom Software	114
Exporting Inventory Information to a File	115
Configuring Monitoring for a Machine	116
Removing a Machine from a Domain	118
Turning Auto Refresh On or Off	119
New Application Domain Panel	120
Installed Software Dialog	122
Machines Dialog	125
View Machine Dialog	127

General Tab	127
Processes Tab	128
Configure Monitoring Tab	128
Add Event Dialog	130
General Pane	130
Alert Pane	131
Required Configuration for sending an email	131
Email Pane	132
Command Pane	133
Chapter 7 Creating and Deploying Applications	135
Application Management Overview	136
Creating an Application	137
Application Creation Choices	138
Application Types	139
Deleting an Application	140
Deploying an Application	141
Deployment Choices	142
Reverting to a Previously Deployed Application	144
Undeploying an Application	145
Undeploy Dialog	145
Deploying an Application Using Dynamic Encryption Key	146
Managing Folders	147
Moving an Application to a Folder	148
Viewing Application Deployment History	149
Purging Application Revisions	150
Upgrading an Application	151
Application Management Dialog	154
All Applications Dialog	155
New Application Configuration Dialog	156
Application Archive Pane	156
Application Parameters Pane	156
Services Pane	157
Deploy Configuration Dialog	159
View History Dialog	161
Chapter 8 Setting Deployment Options	162
Configuration Console Overview	163
Changing Global Variables at Deployment	164

Setting Application Repository Instance Options	167
Enabling a Process or Service to Run on Other Machines	168
Adding a Custom Rulebase to a Process or Service	170
To Add a Custom Rulebase to an Application	170
How to Create a Custom Rulebase	172
Adding an Event to a Process or Service	174
To Add an Event to a Service	174
To set Events for a Process	176
Setting Fault Tolerant Options for a Process	178
To Set Fault Tolerant Options	178
Changing the Checkpoint Data Repository for a Process	179
Configuring Fault-Tolerant Engines	179
Peer or Master and Secondary Relationships	181
Failover and Checkpoint Data	182
Process Starters and Fault-Tolerance	182
Changing Global Variables for a Process or Service	184
Configuring Storage for TIBCO BusinessWorks Processes	185
Specifying a Database for Storage	185
Database Table Names	186
Manually Creating Database Tables	186
Changing TIBCO BusinessWorks Process Configuration Properties	188
To Change Process Configuration Properties	188
Controlling Execution of TIBCO BusinessWorks Services	188
Specifying the Maximum Number of Concurrently Active Processes	189
Specifying Maximum Number of Concurrent Processes in Memory	190
Keeping Services in Memory	190
Effects of Setting the Configuration Fields	190
Changing Server Settings	193
Setting Graceful Shutdown Properties for a Process Engine	194
Specifying HTTP Servlet Authentication Information	195
Application Management Configuration Dialog	196
Configuration Builder Pane	196
Deployed Configuration Pane	198
Edit Application Configuration Dialog	199
General Tab	199
Advanced Tab	199
Edit Service Configuration Dialog	204
General Tab	204
Monitoring Tab	206
Advanced Tab	207
Edit Service Instance Dialog	209

General Tab	209
Server Settings Tab	209
Graceful Shutdown Tab	211
Chapter 9 Managing and Monitoring Process Engines and Service Instances	212
Process Engines and Service Instances Overview	213
Starting or Stopping a Service Instance or Process Engine	215
Viewing Log File Information	217
Editing Process Engine Properties	219
Viewing the TIBCO Administrator Audit Log	220
All Service Instances Dialog	221
View Service Instance Dialog	223
BW Processes	223
General Tab	224
Graceful Shutdown Tab	225
Tracing Tab	226
Security Tab	226
View Service Instance: TIBCO Administrator Dialog	227
General Tab	227
Audit Log Tab	227
Security Tab	227
Plug-Ins Tab	228
Chapter 10 Deploying, Starting and Monitoring an Adapter	229
Prerequisites	230
Required Software	230
See Also	230
Opening the Project in TIBCO Designer	231
Modifying the Adapter Service and Building the Archive	232
Modifying the Adapter Service	232
Building the Archive	234
Creating the Application in TIBCO Administrator	236
Configuring the Application	238
Setting the Variable for the Service	238
Setting Application Options	239
Deploying the Application	241
Starting the Application	243
Monitoring the Application	244
Stopping the Service Instance	246

Chapter 11 Deploying, Starting and Monitoring a TIBCO BusinessWorks Service.	247
Overview	248
Overview of Example Process	248
Prerequisites.	249
Starting TIBCO Designer and Saving Your Project	250
Creating the FileTest Process	251
Testing the FileTest Process	257
Creating the Enterprise Archive File	259
Creating and Deploying the Application in TIBCO Administrator.	260
Starting the Application	261
Monitoring the Application	262
Viewing Default Monitoring Information	262
Specifying a Custom Alert	264
Stopping the Application	266
TIBCO Product Documentation and Support Services	267
How to Access TIBCO Documentation.	267
How to Contact TIBCO Support	267
How to Join TIBCO Community	268
Legal and Third-Party Notices	269

Figures

Figure 1	TIBCO Administrator GUI.	16
Figure 2	Installed Software.	17
Figure 3	Configuration Builder	18
Figure 4	TIBCO Administration Domain	20
Figure 5	Service Instances.	26
Figure 6	Inheriting Group Membership.	45
Figure 7	A Tool Tip Displayed on Mouse Moved Over Role Name	46
Figure 8	Create an Administrator Defined Role	50
Figure 9	Select Groups to Synchronize	52
Figure 10	Add Group	52
Figure 11	Include Descendents Check Box	53
Figure 12	User Management	58
Figure 13	Edit Security	60
Figure 14	Add Roles Dialog	65
Figure 15	Edit Role Dialog	67
Figure 16	Delete a User	71
Figure 17	Manage Password Policy	76
Figure 18	Remove a Role from Its Parent Role	83
Figure 19	Assigning Permissions to Roles.	84
Figure 20	Security Overview	95
Figure 21	TIBCO Administrator Folder.	96
Figure 22	Grant Super User Access	98
Figure 23	Granting Access to an Object	99
Figure 24	Object Locked	101
Figure 25	Message On Breaking Lock	101
Figure 26	TimerProcessArchive Application Inheriting Security Settings.	103
Figure 27	Authorization Inheritance	104
Figure 28	Setting DATA Access Permission	109

Figure 29	Assigning an Application to an Application Domain	110
Figure 30	Advanced Tab	110
Figure 31	Disable Installed Software	113
Figure 32	Auto Refresh Icon	119
Figure 33	Application Management	136
Figure 34	Creating an Application	138
Figure 35	Uploading EAR File	139
Figure 36	Deploying an Application	141
Figure 37	Reverting to a Previously Deployed Application	144
Figure 38	Upgrading a Software	152
Figure 39	Configuration Console	163
Figure 40	Changing Application Properties	166
Figure 41	Enabling a Process or Service to Run on Other Machines	168
Figure 42	Select One or More Machine	169
Figure 43	Adding a Custom Rulebase to an Application	171
Figure 44	Results Displayed When Conditions Specified in the Rulebase Occur	172
Figure 45	Add Event Panel for a Process Archive	175
Figure 46	View Services Instances	176
Figure 47	Normal operation: master processing while secondary stands by	180
Figure 48	Fault-tolerant failover	181
Figure 49	Specifying a Database for Storage	185
Figure 50	All Service Instances	213
Figure 51	TIBCO Designer	231
Figure 52	Modifying the Adapter Service	233
Figure 53	Global Variables Tab	233
Figure 54	Building the Archive	234
Figure 55	Select a Resource	235
Figure 56	New Application Configuration	237
Figure 57	Setting Variable for the Service	239
Figure 58	Setting Application Option	240
Figure 59	Deploying the Application	241
Figure 60	Enter Description for Deployment	241

Figure 61	Configuration Console	242
Figure 62	Starting the Application	243
Figure 63	Monitoring the Application	244
Figure 64	Stopping the Service Instance	246
Figure 65	Save Project Dialog	250
Figure 66	Adding Activities to the Process.	252
Figure 67	Enter the Details in Configuration Tab	253
Figure 68	Activity Icons	253
Figure 69	Mapping Data Flow between Activities.	254
Figure 70	Starting the Application	261
Figure 71	Viewing Default Monitoring Information	263
Figure 72	Configure Tracing.	263
Figure 73	Specifying a Custom Alert	264
Figure 74	Add Event.	265

Tables

Table 1 Command Line Utilities 24

Table 2 Icons Associated with Users in a Domain Using LDAP Directory 37

Table 3 Password Policy Elements 76

Table 4 Status Icons 125

Table 5 Conversion of Domain Names 186

Table 6 Effects of various configuration settings. 190

Chapter 1 **Introduction**

This chapter introduces the TIBCO Administrator™ GUI, the administration domain and other utility programs available in the TIBCO Administrator installation.

Topics

- [Overview, page 16](#)
- [TIBCO Administration Domain, page 19](#)
- [Utilities, Plug-ins and Modules, page 23](#)

Overview

TIBCO Administrator includes two main components, the administration server and TIBCO Administrator GUI. The administration server manages resources in an administration domain. The TIBCO Administration GUI provides a web browser interface, allowing you to configure users and applications, deploy applications, and monitor processes and machines in an administration domain.


This manual explains the TIBCO Administrator GUI interface. See the *TIBCO Administrator Server Configuration Guide* for information about the administration server and administration domain.

The following modules are provided in the TIBCO Administrator GUI:

- [User Management](#)
- [Resource Management](#)
- [Application Management](#)

The next diagram shows the TIBCO Administrator GUI. Clicking the View icon in the top frame allows you to view the configuration panel only, or navigation and configuration panel.

Figure 1 TIBCO Administrator GUI



User Management

This module allows you to create users and roles and assign them access rights to resources available in the administration domain.

An administration domain can be integrated with an LDAP directory server where users and groups are defined in an LDAP directory and imported into TIBCO Administrator. See [Chapter 3, Integrating TIBCO Administrator with an LDAP Directory Server, on page 36](#) for details.

- **Users**—You can create users and assign them into roles, which allow easier administration. You can also set read, write or administer permissions so a

user can directly access TIBCO Administrator modules and consoles, and domain and application repositories.

- **Roles**—Access control is easier when roles are used. You can assign multiple users into roles and then assign permissions for the role, rather than each individual user.
- **Security**—Each resource in an administration domain can have users or roles assigned to it. The security level setting determines who can access the resource and at what access level.

Resource Management

This module allows you to create application domains, get information about installed TIBCO software on each domain machine, view the status of each domain machine, and configure monitoring rules and events that can trigger other actions such as sending email or running a command.

- **Application Domains**—If your TIBCO application supports this feature, you can create multiple application domains and assign applications to use them. An application domain allows you to specify that application data be written to a repository that is separate from the repository used by the administration domain. This is useful, for example, if an application needs to use a local database rather than that used by the administration domain.
- **Installed Software**—You can view TIBCO applications that are running on each machine in the domain and enable or disable applications.

Figure 2 Installed Software

Software	Install Date	Component Path	Machine	IP Address	OS/Version	Enabled
TIBCO ActiveMatrix Adapter for Files for Unix/Win 7.0.0	28-Jun-18 18:34:31 IST	C:/tibco/adapter/adfiles/7.0	ravkumar-p51	10.97.243.10	Windows 10 10.0	Enabled
TIBCO ActiveMatrix BusinessWorks 5.14.0	28-Jun-18 18:09:48 IST	C:/tibco/bw/5.14	ravkumar-p51	10.97.243.10	Windows 10 10.0	Enabled
TIBCO Administrator 5.11.0	28-Jun-18 18:06:32 IST	C:/tibco/administrator/5.11	ravkumar-p51	10.97.243.10	Windows 10 10.0	Enabled

- **Machines**—Each administration domain contains one or more machines. You add a machine to an administration domain using the TIBCO Domain Utility.

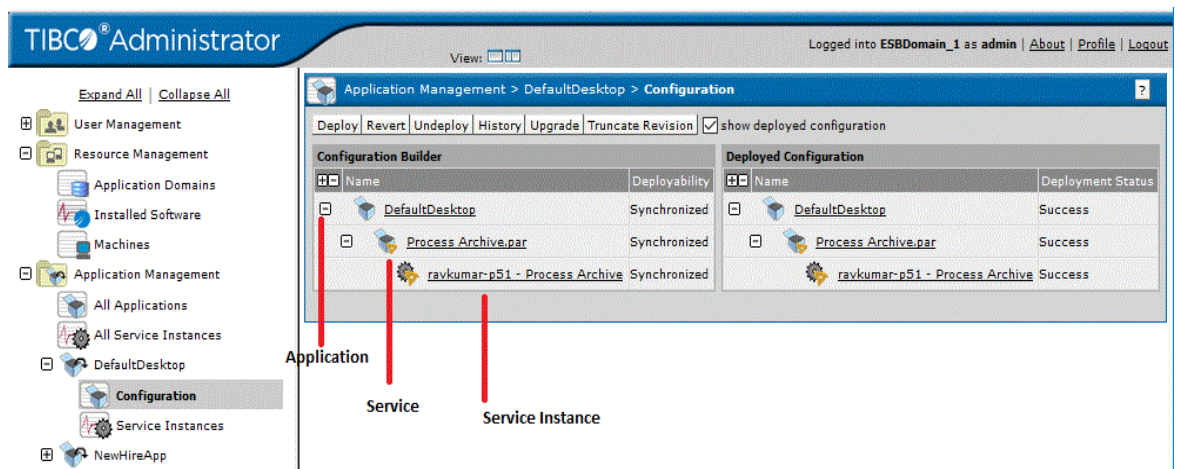
Application Management

This module allows you to upload an application's Enterprise Archive (EAR) file and optionally change options and global variables that were set for the application when it was configured. You can also define monitoring rules for each application. You then deploy the application and start (or stop) it.

The next diagram shows the TIBCO Administrator GUI with the application management module displaying its contained consoles. An application is selected and the configuration builder in the right panel displays the application, its service and service instance.


- Clicking an application name displays a panel where you can change global variable values set for the application when it was configured in TIBCO Designer. If TIBCO Rendezvous is set as the administration domain transport, you can specify the transport the client application will use when communicating with the administration server. You can also specify that the application's repository be sent to the target machine, which allows the application to run independently of the administration server.
- Clicking the service name displays a panel where you can set monitoring options and other properties.
- Clicking the service instance name displays a panel where you can set logging options, whether to run as the instance as a service on Microsoft Windows, and shutdown options.

Figure 3 Configuration Builder



TIBCO Administration Domain

When you create a new TIBCO Silver BusinessWorks instance group in TIBCO Silver Center, it automatically creates an Administration domain. The Administration domain name is the Name field value as shown next.



The screenshot shows the 'Introduction' window of the TIBCO Silver Center. The window has a title bar that says 'TIBCO Silver™ Center'. Inside, the 'Introduction' section contains the following fields and controls:

- Name:** A text field containing 'SilverBusinessWorksDomain'.
- Product Group:** A dropdown menu with 'SilverBW' selected.
- Software Version:** A dropdown menu with 'Silver Business Works 1.0.0.009' selected.
- Network Domain:** A text field containing 'tibco.com'.
- Use VPC:** An unchecked checkbox.

At the bottom of the window are three buttons: 'Cancel', 'Go Back', and 'Continue'.

TIBCO Administrator automatically starts when the master instance of the Instance Group starts.

If you add a new instance to the instance group (by clicking the button to add more instances), the machines will be added automatically to the domain.

When you start an instance, TIBCO Runtime Agent and TIBCO Hawk Agent are started. On each machine, an instance is installed and registered for deployment for the TIBCO BusinessWorks engine.

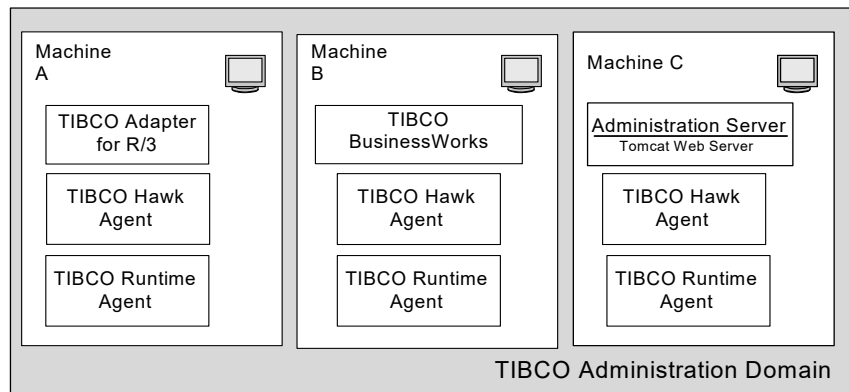


When you remove an instance from an instance domain, the machine is no longer available, but is still registered to the domain. To remove the machine completely you must go to the machine management UI and remove it. See [Removing a Machine from a Domain on page 118](#) for details.

After you install TIBCO Administrator, the installer starts TIBCO Domain Utility to create the initial administration domain. An *administration domain* is a collection of users, machines, and services. Each domain is managed by an administration server, which is assisted by a TIBCO Hawk agent running on each machine in the domain. You can create multiple administration domains on the same machine and each domain must have a unique administration server associated with it.

For example, the next diagram shows an administration domain with three machines. Each machine has TIBCO Runtime Agent installed and a TIBCO Hawk agent running. One machine has an adapter installed, another TIBCO BusinessWorks and the other machine has the administration server installed. The browser-based TIBCO Administrator GUI can be run from any machine, including machines not in the domain.

Figure 4 TIBCO Administration Domain



Administration Server

An administration server manages resources in an administration domain. Each domain has its own administration server. A domain can use TIBCO Rendezvous or TIBCO Enterprise Message Service as the transport for handling domain communication. The transport used affects how data is stored and how applications are deployed. The server's main responsibilities are to:

- Manage data storage for the domain. Domain storage options are specified when creating the domain.

For TIBCO Rendezvous administration domains, domain storage can be in a file-based repository, or in a database repository. For TIBCO Enterprise Message Service administration domains, domain storage must be in a database.

- Manage transport options for applications. Transport options include rv (Rendezvous) http, and local.

For TIBCO Rendezvous administration domains all transport options are available. The default application transport option is set when creating the domain and can be modified per application when deploying the application.

For TIBCO Enterprise Message Service administration domains, applications are always deployed using the local transport option, which sends data to the application's target machine and allows the application to run independently of the administration server.

- Enforce security for the domain.

Only authorized users are allowed access to applications during development. TIBCO Administrator supports both authentication and authorization of users that require read access or write access to applications.

For TIBCO Rendezvous administration domains only, the administration server can:

- Provide load balancing.

Each domain can be configured with one primary server and multiple secondary servers. The primary server allows read and write operations, while secondary servers only support read operations. Only one primary server can be running for each administration domain. Multiple secondary servers are allowed, but each must be on a different machine than the primary server. Secondary servers are defined using TIBCO Domain Utility.



TIBCO Domain Utility does not support promotion of secondary servers to primary servers.

- Provide failure recovery.

You can use a load-balanced administration servers for failure recovery, where secondary administration servers continue serving requests even with the primary administration server down.

See the *TIBCO Administrator Server Configuration Guide* for more information.

Tomcat Web Server

The Tomcat web server manages basic communications and makes the TIBCO Administrator GUI available. The ports used by the Tomcat web server can be configured when using TIBCO Domain Utility to create the domain, or later changed if required.

TIBCO Hawk Agent

The TIBCO Hawk agent is an autonomous process that monitors applications and systems activity. Each administration server has a corresponding TIBCO Hawk agent. The Hawk agent monitors activity by processing rulebases, which hold the logic that determines how monitoring and management will take place. The Hawk agent also builds local properties files and executable files for deployed applications, and creates NT services for applications.

A Hawk agent runs on the machine that hosts the administration server and on each machine that is part of the administration domain.

TIBCO Runtime Agent

TIBCO Runtime Agent provides the runtime environment required for TIBCO applications. It includes many software tools including TIBCO Domain Utility, TIBCO Designer, TIBCO Rendezvous as well as libraries used by many TIBCO applications. Command line utilities used by TIBCO Administrator are included in the TIBCO Runtime Agent installation.

Utilities, Plug-ins and Modules

TIBCO Administrator includes several utilities, plug-ins and modules that are introduced in this section. On the master TIBCO Silver BusinessWorks, TIBCO_HOME is set to the following value: /opt/tibco/.

TIBCO Domain Utility

The TIBCO Domain Utility is part of the TIBCO Runtime Agent installation. The utility is launched directly after installing the administration server and is used to create the initial administrator domain. You can start the utility at any time to manage domain machines, domain configurations, administration server settings, upgrade administration domains and register a TIBCO Enterprise Message Service server. A registered server can be configured using the TIBCO Enterprise Message Service Plug-in that is described later in this section.

TIBCO Domain Utility is documented in the *TIBCO Runtime Agent Domain Utility User's Guide*. The guide is part of the TIBCO Runtime Agent documentation set and is available in the `TIBCO_HOME/tra/<version>/doc` directory.

Scripting Deployment Utility

Three scripting utilities are available, buildear, ImportDomainSecurity and AppManage. The utilities allow you to perform the functions available in the GUI applications from the command line. This allows you to create and execute scripts for repetitive tasks.

The utilities are installed as part of the TIBCO Runtime Agent installation and are documented in the *TIBCO Runtime Agent Scripting Deployment User's Guide*. The guide is part of the TIBCO Runtime Agent documentation set and is available in the `TIBCO_HOME/tra/<version>/doc` directory.

The buildear utility builds an enterprise archive file based on an archive defined in a TIBCO Designer project. You provide the location of an archive resource in TIBCO Designer, location of the TIBCO Designer project and location of the output enterprise archive file to the buildear utility. You can optionally save an archive version to the project when building the enterprise archive file.

The AppManage utility creates an XML based deployment configuration file where deployment options can be defined. The utility also uploads the deployment file and enterprise archive file into an administration domain. It can be used to:

- Create a deployment configuration file based on information in an enterprise archive file, or from an application already configured in the TIBCO Administrator GUI.
- Upload an enterprise archive file to an administration domain without specifying deployment configuration options. After the file is imported, it is ready to be configured and deployed using the TIBCO Administrator GUI.
- Upload an enterprise archive file and a deployment configuration file into an administration domain in one operation. The application is uploaded with its deployment options set, but is not deployed.
- Upload an enterprise archive file and a deployment configuration file and deploy the application in one operation. Using this method, you can quickly deploy your applications in multiple domains.
- Export all application archives and deployment configuration files within a domain, so they can be batch deployed, undeployed or deleted in another domain.
- Undeploy a deployed application.
- Delete an application from an administration domain. If the application is deployed, you can undeploy it and delete it in one operation.
- Start an application's process engines or service instances
- Stop a successfully deployed process engine or service instance.
- Move an application's configuration storage between server and local.
- Change the domain default for application storage.

Command Line Utilities

The following command line utilities are documented in the *TIBCO Administrator Server Configuration Guide*.

Table 1 Command Line Utilities

Command	Description
CorpRoleSynchronizer	Syncs an administration domain with its associated LDAP directory.
CorpUserSynchronizer	Pre loads user objects into an application.
DeleteInvalidUsers	Removes the users from an LDAP domain if the users do not exist in the associated LDAP directory.

Table 1 Command Line Utilities

Command	Description
ExportDomainSecurity	Exports security data from a domain. This includes users and roles.
ImportDomainSecurity	Imports security data into a domain.
MoveMachine	Moves an administration domain from one machine to another.
RedeployAllAppsForUser	Updates password changes to all deployed applications in a given domain.
RepoConvert	Converts back and forth between single-file (.dat) and multi-file projects.
RepoCreateInstance	Creates a new project.
RepoDelete	Deletes nodes from a domain repository.
RepoDeleteInstance	Deletes a repository instance.
RepoDiff	Logically compares two projects and outputs the differences.
RepoExport	Exports all parts of a project to an XML file.
RepoImport	Loads data from a file created with RepoExport into a project.
RepoListInstances	Lists currently available projects found at the location specified by a URL. The URL can either be local or remote.
RepoPing	Checks whether a file based administration server is running and communicating.
RepoRename	Renames a node in a local single-file or server-based project. Do not use this command if you have saved the project in multi-file format.
AppStatusCheck	This utility is used to query status of all deployed applications in a domain.

TIBCO Enterprise Message Service Plug-in

The TIBCO Enterprise Message Service plug-in (EMS plug-in) allows you to use TIBCO Administrator to configure a TIBCO Enterprise Message Service server. The EMS plug-in is included in your TIBCO Administrator installation.

The EMS plug-in includes online help files that describe each parameter that can be set from within the TIBCO Administrator GUI. You should be familiar with the TIBCO Enterprise Message Service documentation, which explains full use of the server.

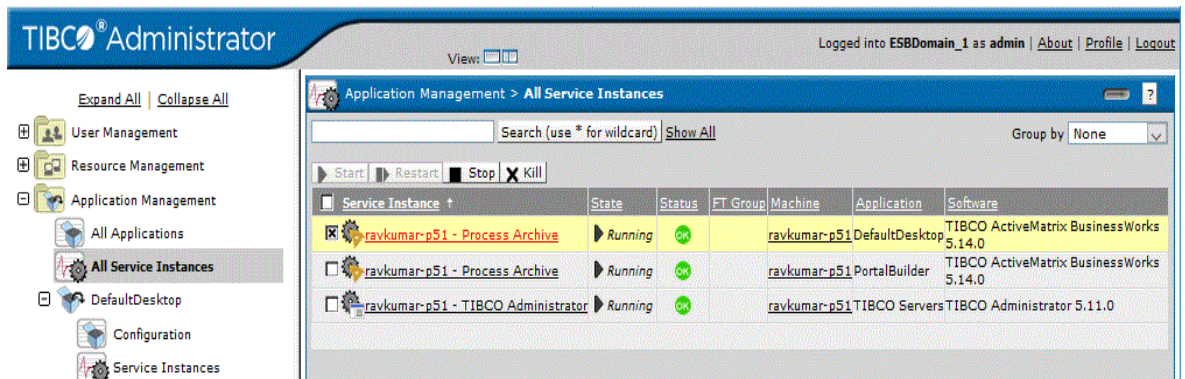
You must first install the TIBCO Enterprise Message Service server, and then run TIBCO Domain Utility to add the EMS plug-in to your administration domain. After the EMS plug-in is installed, you access it from the TIBCO Administrator GUI by clicking **Application Management > All Services Instances > machine-name - EMS port**.



After you install TIBCO Enterprise Message Service, you must stop the TIBCO EMS Server windows service. You should also make sure that no such service (tibemspd.exe) is running before you start the EMS plug-in from the TIBCO Administrator GUI.

The next diagram shows the All Service Instances panel with the EMS server status listed.

Figure 5 Service Instances



The users and roles defined in your domain can be synchronized with the users and groups defined for the TIBCO Enterprise Message Service server. Both users and their passwords become known to the server after synchronization.



The synchronization removes all existing Enterprise Message Service users and groups except the admin user and group. The admin user and admin group are not changed.

If your domain is integrated with an LDAP directory server, synchronization with the users and groups defined for the TIBCO Enterprise Message Service server is not supported. An LDAP directory server does not allow passwords to be exported, and as a result passwords would not be exported on synchronization.

To Set Parameters Using the EMS Plug-in

1. Click **Application Management > All Service Instances**
2. Expand the TIBCO Enterprise Message Service.
3. Click the server instance name.

You can set options for the following:

- Server Parameters
- Queues Parameters
- Topics Parameters
- Durables Parameters
- Connection Factories Parameters
- Connections Parameters
- Producers and Consumers Parameters
- Routes Parameters
- Bridges Parameters
- Transports Parameters
- Transactions Parameters
- Users and Groups Parameters
- JNDI Bindings Parameters

TIBCO BusinessWorks Manual Work Module

TIBCO BusinessWorks provides a simple web interface for viewing manual work tasks assigned to users. This interface is named BW Manual Work in the TIBCO Administrator GUI. This will only be available if TIBCO BusinessWorks or TIBCO InConcert is installed on the TIBCO Administrator server machine.

The BW Manual Work module is documented in *TIBCO BusinessWorks Palette Reference*, which is part of the TIBCO BusinessWorks documentation set.



The TIBCO BusinessWorks Manual Work module does *not* support file-based (non-database) administration domains.

To access the module from the TIBCO Administrator GUI, in the left panel, click **BW Manual Work**.

When you log into the web interface, you can see any task that you have acquired. To log in, you must specify a valid username, password, and TIBCO InConcert server name.

Authorized users can perform the following tasks:

- Export Users or Roles to TIBCO InConcert
- Specify a TIBCO InConcert Proxy Server
- View Task Lists
- Acquire Tasks
- Work with Tasks in the Acquired List
- Work with the Completed List
- Administer Tasks

UDDI Servers Module

The Universal Description, Discovery, and Integration (UDDI) Servers module allows you to define connections to UDDI servers and view the web services contained in the servers. If you have been granted access to publish your own web services, you can also use the UDDI Servers module to publish information about your business and the web services you offer. See the *TIBCO BusinessWorks Administration* guide for information about UDDI.

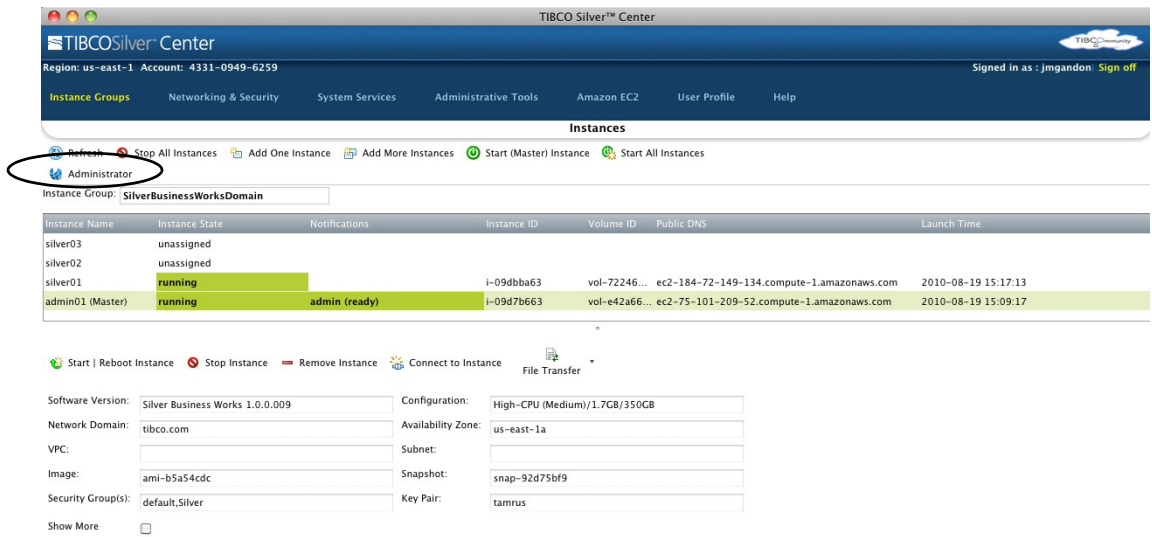
Chapter 2 Starting TIBCO Administrator

This chapter explains how to start and stop the administration server and the TIBCO Administrator GUI.

When you start the master instance of an instance group it automatically starts TIBCO Administrator and the TIBCO Hawk Agent. No action is necessary from the user.

Starting the TIBCO Administrator GUI

To access to the TIBCO Administrator GUI, click the **Administrator** button in Silver Center as shown next.



This will launch your browser with the following URL:
`https://Public-IP-Address/administrator/servlet/tibco_administrator`
Your default browser must be set to one of the following:

- Microsoft Internet Explorer 5.5 or higher
- Mozilla Suite 1.7.1 or higher
- Mozilla Firefox 1.x or higher

Note: Admin is accessible only when the master instance status is Admin (ready).
 Login in administrator with the user password that has been specified during the Instance group creation (see attached screen).

For the first login, you must use the user specified as the domain administrator user with Domain Utility. That user can then assign other users privileges to log in.

Stopping the Administration Server

When you stop the master instance, it will stop TIBCO Administrator

You can also use the appropriate kill command for your system to stop the administrator server.

For a graceful shutdown:

1. On the command line, change to
`TIBCO_HOME/administrator/domain/domain/bin.`
2. Execute: **tibcoadmin_domain -stop.**

Note that the above command syntax initiates an asynchronous shutdown sequence.

Stopping the Administration Server Remotely

You define two parameters when creating an administration domain, Shutdown Port and Shutdown String. You can use these two parameters to shut down the administration server remotely.

1. In a command prompt, execute the following command, where *hostname* is the name of the administration server host and *port* is the shutdown port:

```
telnet hostname port
```

2. Type the shutdown string and press Enter.

You may not observe anything on the command prompt except that the operation exits. However, you can verify on the administration server host that the administration server is indeed shut down.

Topics

- [Starting TIBCO Administrator on Microsoft Windows, page 32](#)
- [Starting TIBCO Administrator on UNIX, page 34](#)
- [Stopping the Administration Server, page 35](#)

Starting TIBCO Administrator on Microsoft Windows

You must start the administration server and TIBCO Hawk agent before starting the TIBCO Administrator GUI.

Starting the Administration Server and TIBCO Hawk Agent

The administration server is normally started immediately after the initial administration domain is created. Two Microsoft Windows services must be running for the server to be available. The services are installed by the installation program and set to start automatically. To start the services the first time, navigate to the Services dialog, find the administration server for your domain and click the **Start** button. Repeat for the TIBCO Hawk agent service.



If the TIBCO Hawk agent is started as a service, mapped drives on the machine are not recognized by deployed services. The workaround is to start the TIBCO Hawk agent from the command line.

Alternatively, to start on the command line:

1. Start the administration server by typing the following into a command-line prompt:

```
% cd TIBCO_HOME/administrator/domain/domain/bin
% tibcoadmin_domain-name.exe
```

2. Start TIBCO Hawk agent by typing:

```
% cd TIBCO_HOME/tra/domain/domain
% hawkagent_domain-name.exe
```

If any value is modified in the `hawkagent_domain.tra` or `tibcoadmin_domain-name.tra` file, then you must update the corresponding service. For `hawkagent`:

```
% cd TIBCO_HOME/tra/domain/domain-name
% hawkagent_domain-name.exe --update
```

For `tibco_domain-name.tra`:

```
% cd TIBCO_HOME/administrator/domain/domain-name/bin
% tibcoadmin_domain-name.exe --update
```

Starting the TIBCO Administrator GUI

You can launch the TIBCO Administrator GUI by entering the appropriate URL into your browser, or you can use the **Start** menu.

Starting from a Web Browser

1. Open a web browser and connect to this URL for the TIBCO Administrator GUI:

`http://host-name:port/administrator/servlet/tibco_administrator`

- *host-name* is the name of the machine on which the administration server has been installed. If this is the same machine you are currently on, you can use **localhost** as the machine name.
- *port* is 8080 by default. If you have used the TIBCO Domain Utility to assign a different port, use that port number instead. If you created multiple domains on one machine, the port is incremented by 10 for each domain. For example, the second domain will use 8090.



You can enter `http://host-name:8080`. This displays a list of domains, the port each domain is using, and the TIBCO software available on that port. You can pick a domain from this list to go to the login screen.

2. Log in. For the first login, this must be the user specified as the domain administrator user with the Domain Utility. That user can then assign other users privileges to log in.

Starting from the Start Menu

To start TIBCO Administrator from the Start menu, your default browser must be set to one of the following:

- Microsoft Internet Explorer 7.x, 8.x
- Mozilla Firefox 3.x

Follow these steps:

1. Select **Start > All Programs>TIBCO>TIBCO Administrator *n.n* > TIBCO Administrator**.
2. Log in. For the first login, this must be the user specified as the domain administrator user with the Domain Utility. That user can then assign other users privileges to log in.



You must enable cookies in the browser for TIBCO Administrator to work.

Starting TIBCO Administrator on UNIX

You must start the administration server and TIBCO Hawk agent before starting the Administrator GUI.

Starting the Administration Server and TIBCO Hawk Agent

The administration server is normally started immediately after the initial administration domain is created.

1. First start the server by typing the following into a command-line prompt:


```
% cd TIBCO_ADMIN_DOMAIN_HOME/domain-name/bin
% tibcoadmin_domain-name
```
2. Then start the TIBCO Hawk agent, which performs the TIBCO Administrator monitoring functions, by typing:


```
% cd TIBCO_TRA_DOMAIN_HOME/domain-name
% hawkagent_domain-name
```

Starting the TIBCO Administrator GUI

1. Open your web browser and connect to the following URL:


```
http://host-name:port/administrator/servlet/tibco_administrator
```

 - *host-name* is the name of the machine on which the administration server has been installed. If this is the same machine you are currently on, you can use `localhost` as the machine name.
 - *port* is 8080 by default. If you have used the TIBCO Domain Utility to assign a different port, use that port number instead. If you have created multiple domains on one machine, the port is incremented by 10 for each domain. For example, the second domain will use 8090.



You can enter `http://host-name:8080` to get a list of domains, the ports they are using, and the TIBCO software available on that port, and then pick a domain from this list to go to the login screen.

2. Log in as the domain administrator user. This user was specified using the Domain Utility. That user can then assign other users privileges to log in.

Stopping the Administration Server

You can stop the administrator server in several ways:

- On all platforms, from the TIBCO Administrator GUI, choose **Application Management > All Services Instances**. Select *machine-name - TIBCO Administrator* and click **Stop Selected**. Note that after stopping the server, the Administrator GUI becomes unavailable.
- If you started the administration server from a command line, you can use Control-C in the command window that launched the server on any platform to stop the server.
- On Microsoft Windows, navigate to the Services panel. Select the administrator server service, and then click the **Stop** button.
- On UNIX, use the appropriate kill command for your system to stop the administrator server.

For a graceful shutdown, follow the steps below:

- a. In the command line, change to
`TIBCO_HOME/administrator/domain/domain/bin.`
- b. Execute: `tibcoadmin_domain stop.`

Note that this command syntax initiates an *asynchronous* shutdown sequence.

Chapter 3

Integrating TIBCO Administrator with an LDAP Directory Server

This chapter explains how to use users and groups defined in an LDAP directory server in an administration domain.

Topics

- [Overview, page 37](#)
- [Managing LDAP Users and Group-synchronized Roles, page 45](#)
- [Adding a Local User to an LDAP Integrated Domain, page 48](#)
- [Adding a Local Role to an LDAP Integrated Domain, page 49](#)
- [Filtering LDAP Users and Groups to Integrate, page 51](#)
- [Pre Loading User Objects, page 55](#)
- [LDAP Synchronization Optimistic Option, page 56](#)

Overview

Most enterprises use an LDAP directory server that provides a basis for authorization for all its enterprise applications. An LDAP directory contains information about users and the groups to which users belong. Groups can also include other groups as child groups. In some cases LDAP directories also contain information about customers and vendors, providing authorization for customer service and supply chain applications.

By integrating an administration domain with an LDAP directory server, TIBCO applications and services can leverage the users and groups from the LDAP directory server. Note that an LDAP directory is referred as Corporate LDAP in various places in TIBCO Domain Utility.

Local users and local roles can be created in an administration domain that is integrated with an LDAP directory server. The following table shows the icons associated with the different users and roles that are available in a domain that uses an LDAP directory.

Table 2 Icons Associated with Users in a Domain Using LDAP Directory





Icon	Description
	<p>Local user</p> <p>A local user is created in the TIBCO Administrator GUI. The user can be assigned to local roles, but cannot be assigned to LDAP group-synchronized roles. Local users are authenticated against credentials stored in the TIBCO authorization domain. No information about a local user is stored in the LDAP directory.</p>
	<p>LDAP user</p> <p>An LDAP user is created in the LDAP directory. An LDAP user can be assigned into local roles in the TIBCO Administrator GUI, but cannot be assigned to other LDAP group-synchronized roles.</p> <p>LDAP users are authenticated against credentials stored in the LDAP directory. An LDAP user cannot be deleted from the LDAP directory using the TIBCO Administrator GUI.</p> <p>Note: LDAP users <i>cannot</i> be authenticated using password digests.</p>
	<p>Local role</p> <p>A local role is created in the TIBCO Administrator GUI. Local users, LDAP users, or both can be placed in local roles. No information about local roles is stored in the LDAP directory.</p>

Table 2 Icons Associated with Users in a Domain Using LDAP Directory

Icon	Description
	<p>LDAP Group-synchronized role</p> <p>An LDAP group-synchronized role is an LDAP group that is created in an LDAP directory. LDAP groups can be placed in local roles, but cannot be placed in other LDAP group-synchronized roles.</p>

Supported LDAP Directory Servers

The following LDAP directory servers are supported:

- Sun ONE Directory Server 5.2
- Sun ONE Directory Server 5.1 with Service Pack 2
- Microsoft Active Directory 2003, 2008, 2008 R2
- Microsoft Active Directory Lightweight Directory Service (LDS) Application Mode (ADAM) 2003, 2008, 2008 R2
- Novell eDirectory 8.6.2, 8.7.3, 8.8.2, 8.8.4
- CA Directory Server r8.1
- OpenLDAP 2.3

Features

The following major features are supported.

Leverage Users

TIBCO applications and services can authenticate users against an LDAP directory and get read access to an LDAP user’s properties.

TIBCO Administrator does *not* create a copy of the user present in an LDAP directory. If TIBCO Administrator (or other TIBCO applications) requires LDAP user properties at runtime, TIBCO Administrator retrieves the properties directly from the LDAP directory and caches the properties in memory with a suitable expiry time. Group membership is also retrieved at runtime and cached in memory with a suitable expiry time.

TIBCO Administrator creates a user entity (an object) in its database in which only the username (or user login id) is copied from the LDAP directory. The object stores specific user profiles for use with TIBCO applications.

Local users can be created in an administration domain that is integrated with an LDAP directory.

Leverage Groups

TIBCO applications and services can use both static and dynamic LDAP groups available in an LDAP directory in the form of LDAP group-synchronized roles.

In addition to LDAP group-synchronized roles, local roles can be created in TIBCO Administrator. Local roles can include LDAP group-synchronized roles as members.

Dynamic Configuration Changes

Configuration changes in an LDAP directory server, such as the addition, deletion or modifications of user and groups do not require TIBCO applications or services to be restarted. This latest information becomes available after the next LDAP synchronization or after the expiry interval is triggered.



If you change LDAP connection parameters, search parameters or synchronization parameters, TIBCO applications and services, including the administration server must be restarted. The parameters are changed using TIBCO Domain Utility.

Connection to an LDAP Directory Server

TIBCO applications and services, including the TIBCO Administrator server can connect to an LDAP directory server using simple (or basic) authentication, or SSL (Secure Sockets Layer) authentication.

- Basic authentication is the simplest security mechanism available to connect to an LDAP directory server. When using basic authentication, the TIBCO Administrator server identifies itself to the LDAP directory server by means of a Bind DN (Distinguished Name) and a password, which are sent in clear text over the network.
- All data exchange between the TIBCO Administrator server and the LDAP directory server, including the Bind DN and password used while creating the connection, can be secured using an SSL connection. See the *TIBCO Administrator Server Configuration Guide* for details.

Support for Server-side Chaining and Client-side Referrals

LDAP directories provide two mechanisms to structure their Directory Information Tree (DIT) in a distributed manner: server-side chaining and client-side referrals. While searching through distributed LDAP directories, a query may need to span and traverse multiple directory servers. An administration domain can be integrated with these distributed LDAP directories.

In the case of server-side chaining, the responsibility to traverse to the chained data lies with the LDAP directory server. TIBCO applications need not do any special processing. An administration domain is configured to connect to the primary LDAP directory in the usual way.

In the case of client-side referrals, the responsibility to traverse other directories is with the client, that is, TIBCO applications or services. While searching a distributed LDAP directory, the referred LDAP directory URL is provided to the client and it traverses to that LDAP directory to collect matching LDAP entries.

Configuration through TIBCO Domain Utility

To connect to an LDAP directory server, TIBCO applications and services need to know the server's bind information. This information is configured using Domain Utility in the domain properties screen. LDAP directory connection information includes bind information for the primary LDAP directory and for all referral LDAP directories. Bind information is validated at configuration time.

Configuration also involves specifying the search parameters for users and groups. This information is optionally validated at configuration time. The search parameter settings can be saved even if validations fail.

Synchronization parameters such as synchronization and expiry intervals are also configured through Domain Utility. Note the following distinctions between synchronization of LDAP groups and LDAP users and the administration domain.

- LDAP groups are synchronized periodically. The synchronization operation creates an LDAP group synchronized role in the administration domain for each LDAP group.
 - Membership is not imported the administration domain, but is lazy loaded on demand and cached with an expiry period (default is 10 minutes).
 - The default synchronization interval is 12 hours

- LDAP users are not synchronized or imported periodically.
 - LDAP users are available to any TIBCO application as synchronized users. User names are queried directly from the LDAP directory.
 - Empty user profiles are created on demand for synchronized users (typically when the user logs in for the first time). You will see a corresponding entry in the audit.log for this event.
 - LDAP properties of synchronized users are not imported into the administration domain, but are lazy loaded on demand with an expiry period (default is 10 minutes). User profiles in the domain are meant to store only local TIBCO application specific properties such as subscriptions.
 - The CorpUserSynchronizer command line utility is provided to create empty profiles of all synchronized users up front. See [Pre Loading User Objects on page 55](#) for details.

NetBIOS Domain-based Names

When using Microsoft Active Directory as the LDAP directory server, an option is available to use NetBIOS domain-based names. This option adds NetBIOS domain names as a prefix to user names and group-synchronized roles. This allows LDAP directories that contain users or roles of the same name across different domains to be used in TIBCO Administrator.

Log Files

If an LDAP directory server invocation error occurs, the complete error message displays on the console and is also written to the `Administrator.log` file so that you can manually recover and process the message. The log file is written in the `TIBCO_HOME/tra/domain/domain/logs` folder.

Tool Tips

In the TIBCO Administrator GUI, tool tips are displayed for a role when a mouse is moved over the role's name. A tool tip displays a role's name, description, and paths to this role, based on role hierarchy. It also displays the LDAP DN (Distinguished Name) of the corresponding group, for a group-synchronized role.

Limitations

The following are not supported when integrating an administration domain with an LDAP directory server.

No Writes to an LDAP Directory Server

TIBCO applications and services do not write to an LDAP directory server. Changing an LDAP group or its membership information from the TIBCO Administrator GUI is not supported.

Limited to Users and Groups

Information loaded from an LDAP directory server is limited to users, groups and its membership information. TIBCO applications and services do not synchronize, retrieve or use any other information from an LDAP directory.

No Advanced Authentication

Users of TIBCO applications and services are authenticated against an LDAP directory using basic authentication. Other aspects of authentication such as prompting for password changes or displaying password expiry notices in TIBCO applications are not supported.

Restrictions on Users and Roles

LDAP users and group-synchronized roles cannot be deleted or renamed in the TIBCO Administrator GUI for an LDAP integrated domain. Local users and roles created using the TIBCO Administrator GUI are allowed to be deleted or renamed.

Any role, including group-synchronized roles, can be added as a member or child of a locally created role in the TIBCO Administrator GUI. However you cannot add a local role as member or child of a group-synchronized role.

LDAP Directory Server Must be Running

If an LDAP directory server is down TIBCO applications and services will be unable to get information about LDAP users and groups. This means the TIBCO application is dependent on the LDAP directory server. In most cases an LDAP directory server also serves as the authentication source and needs to be up and running in order for any user to login to a TIBCO application or service.



When the administration server (or other TIBCO applications that connect to an administration domain) starts and you find an error in the server's log file that indicates the LDAP directory server is down, the administration server (and other TIBCO applications) must be restarted after starting the LDAP directory server. If, however, the LDAP directory server goes down later in the session and is brought back, there is no need to restart the server (or other TIBCO applications).

Renaming Groups in an LDAP Directory

Though it is rarely used, an LDAP directory server allows group names to be renamed. A renamed group name can cause problems if a TIBCO application or service determines that a renamed group is a new group and that the old group has been deleted. This results in deleting and creating a new LDAP group-synchronized role for the renamed LDAP group. This can affect the access control list of resources that referred to the original LDAP group-synchronized role. It also affects other locally created roles that included the original LDAP group-synchronized roles as its member or child.



If your administration domain is integrated with Microsoft Active Directory, group renames will be detected during the LDAP synchronization cycle.

Searching and Active Directory Server

If your administration domain integrates with an LDAP directory for users and groups, it requires TIBCO Runtime Agent based applications (such as TIBCO Administrator) to search the LDAP directory. The search can include activities such as synchronizing roles with corporate groups, synchronizing with corporate users, retrieving corporate group membership and searching for users.

If Active Directory is used, it forces a limit of 1000 entries on a regular search. However, a paged search feature (Virtual List View) can be used to retrieve more than 1000 entries. The paged search feature works with Active Directory 2003 only (Service Pack 1 for Windows Server 2003 must be installed). The search limit is due to a defect in Windows Server 2003 that throws an "Unavailable Critical Extension" error message. The following link provides details about this defect:

<http://support.microsoft.com/default.aspx?scid=kb;en-us;886683>

Using the paged search feature, up to 10,000 entries can be returned. If the search involves more than 10,000 entries, the search must be broken into multiple sets of smaller queries. This can be achieved by specifying multiple sets of search parameters for users and groups under LDAP Settings in TIBCO Domain Utility.

For previous versions of Active Directory the search is limited to 1,000 entries. You can either raise this limit on your LDAP directory server using the `ntdsutil` utility that is part of Active Directory server, or specify multiple search parameters with smaller queries, as described in the last paragraph in this section.

If the page size limit, referred as `MaxPageSize` in the Active Directory installation, is configured with a value other than the default value (other than 1000), an additional step must be performed for this feature to work correctly. The value can be viewed using the `ntdsutil` utility. The following parameter must be set in the `AuthorizationDomain.properties` file. The parameter must be set to the actual value of `MaxPageSize`. For example:

CorpLdapMaxPageSize=2000

You should also check the maximum value range limit, referred to as `MaxValRange`, in the Active Directory installation. This value can be viewed using the `ntdsutil` utility. This value affects the search that retrieves membership of a Corporate Group. If this limit is configured with a non default value (other than 1000), the following parameter must be set in the `AuthorizationDomain.properties` file. The parameter must be set to the actual value of `MaxValRange`. For example:

CorpLdapMaxValRange=2000

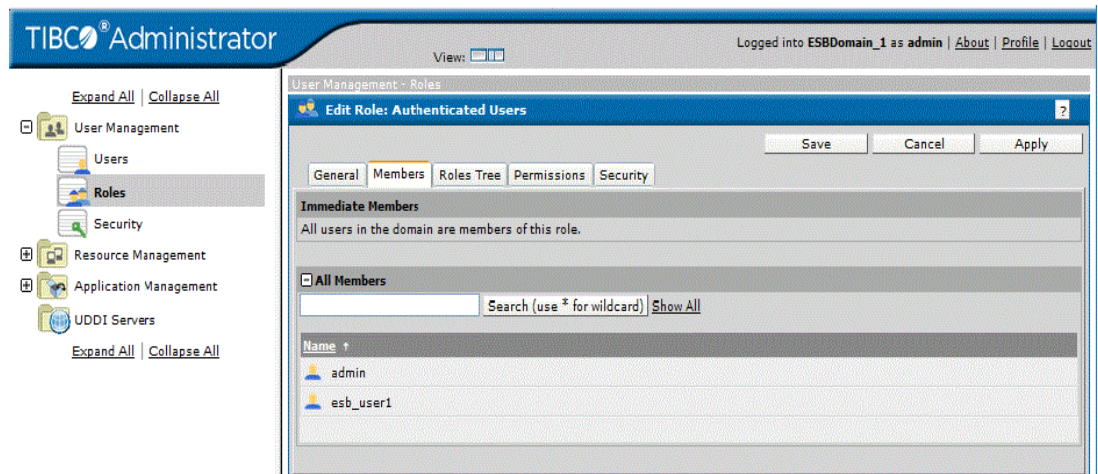
Managing LDAP Users and Group-synchronized Roles

The LDAP users that display in the TIBCO Administrator GUI are defined in an LDAP directory. You cannot create or delete LDAP users in the TIBCO Administrator GUI. LDAP users can be assigned to roles created in TIBCO Administrator, not to LDAP group-synchronized roles synchronized from an LDAP directory server.

An LDAP directory contains users arranged in groups that normally represent the corporate hierarchy. A group has a membership list that contains users and other groups known as child groups. Each child group can have its own membership list that could contain yet other child groups and this leads to a group hierarchy.

Group membership is inherited up the group hierarchy. This means that the members of a child group are implicitly considered to be the members of its parent group. The complete membership of a group is computed by including the members of its child groups. It is also possible for a child group to be a member of more than one parent group. An LDAP directory server does not check for cyclic hierarchies, and thus cyclic hierarchies may exist.

Figure 6 Inheriting Group Membership



The root role, Authenticated Users, is available in addition to LDAP group-synchronized roles. The role hierarchy in the TIBCO Administrator GUI mirrors the LDAP group hierarchy. The LDAP group-synchronized roles that correspond to the top-level group, that is, the groups that do not have a parent in the LDAP directory, are created as child roles of Authenticated Users.

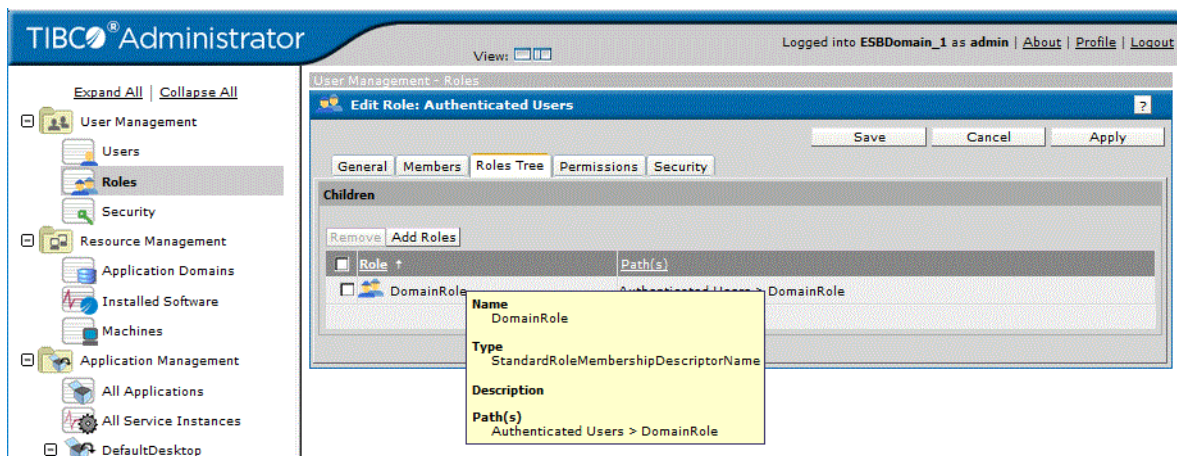
Group-synchronized roles are named using the Relative Distinguished Name (RDN) and not its Distinguished Name (DN).

An LDAP directory can contain two groups with same RDN in different parts of its object tree. However these groups will have a different DN, which uniquely identifies a group. For example the following groups (specified by their DN) have same RDN of Partners:

```
cn=Partners,dc=na,dc=tibco,dc=com
cn=Partners,ou=groups,dc=na,dc=tibco,dc=com
```

In the case where an LDAP directory contains groups with same RDN but different DN, synchronized roles are distinguished either by their relative location in the role hierarchy or by checking the DN for the corresponding group under the tool tip for that role. A tool tip displays when mouse is hovered over the role name in the TIBCO Administrator GUI.

Figure 7 A Tool Tip Displayed on Mouse Moved Over Role Name



When using TIBCO Domain Utility to create or modify an administration domain that is integrated with an LDAP directory, you can choose to automatically create a role for each LDAP group. You can also choose not to create these roles automatically. See the next sections for details.

Automatically Creating a Role for Each LDAP Group

When the **Automatically create Roles for each Corporate Group** feature is selected, a corresponding role is created for each group found in an LDAP directory server. These roles have the same name as their corresponding LDAP group, and the membership of these roles is directly governed by the membership of the LDAP group that it is synchronized with. The membership in this case is fixed and cannot be modified. These roles are referred as LDAP roles or LDAP group-synchronized roles.

The synchronization process used to create LDAP group-synchronized roles is periodically executed in background within a TIBCO Administrator service. Only one primary instance of a TIBCO Administrator service runs this process (in a fault tolerant mode).

Choosing Specific LDAP Groups to Synchronize

When using TIBCO Domain Utility to create an administration domain that uses an LDAP directory, if the **Automatically create Roles for each Corporate Group** feature is not selected, no LDAP group-synchronized roles are created in the TIBCO Administrator GUI for LDAP groups. Instead, each LDAP user is assigned to the root role, **Authenticated Users**, in the TIBCO Administrator GUI.

Even though no LDAP group-synchronized roles are created automatically for each LDAP group, you can still manually synchronize using the Synchronize button in the Select: LDAP Groups screen in the TIBCO Administrator GUI. See [Selecting LDAP Groups to Synchronize in TIBCO Administrator, page 51](#) for more information.

Adding a Local User to an LDAP Integrated Domain

A local user can be created in an administration domain that is integrated with an LDAP directory. The user can be assigned to local roles, but cannot be assigned to LDAP group-synchronized roles. Local users are authenticated against credentials stored in the TIBCO authorization domain. No information about a local user is stored in the LDAP directory.

A local user name can conflict with a LDAP user name when:

- an LDAP user with same name as a local user is later created in the LDAP directory.
- a local user is created with the same name as an LDAP user already defined in the LDAP directory.

In both these cases, the local user takes precedence over the LDAP user. If the local user is deleted, the same-named LDAP user will automatically be exposed.

To Add a Local User

1. Select **User Management > Users** in the left panel.
2. Click **New User**.
3. Supply the user name and click **OK**.
4. Click **set** to provide a password for the user. After entering the password, click **OK**.
5. Click **Save**.

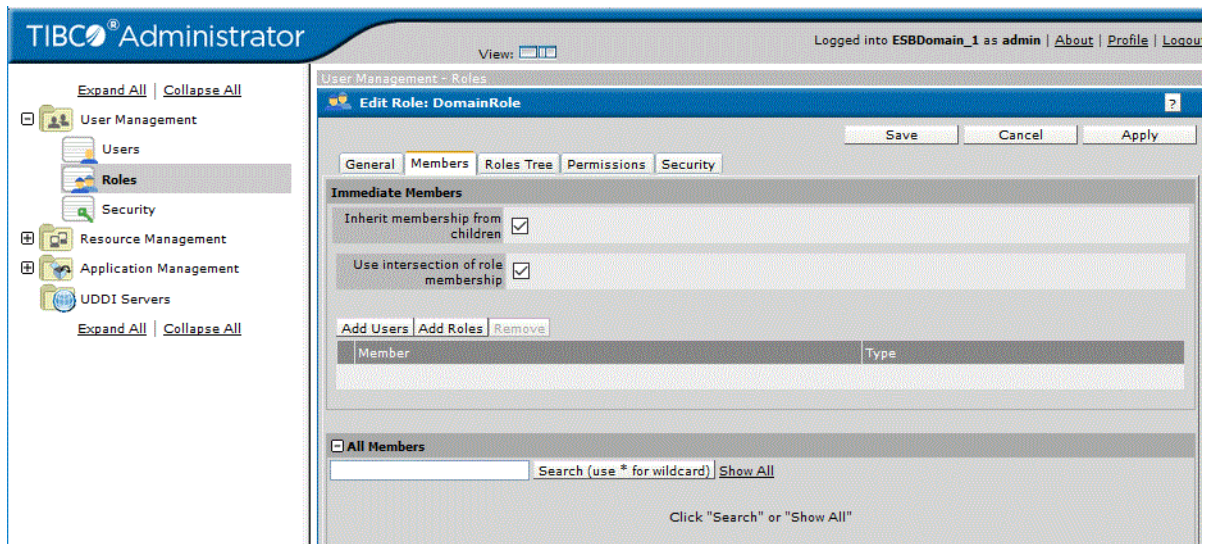
Adding a Local Role to an LDAP Integrated Domain

You can add local roles to an administration domain that is integrated with an LDAP directory. A local role can contain LDAP users and other roles (local or LDAP group-synchronized roles) as its members.

When you add another role to a local role, by default, all members of the added role will be members of the local role. However, if you select the check box **Use Intersection of Role Membership**, only members who are in every added role become members of the local role. Combining roles can be useful, for example, to assign authorization to all members (or the intersection).

To Create an Administrator Defined Role

1. In TIBCO Administrator, in the left panel under User Management, select **Roles**.
2. Select the check box next to **Authenticated Users**.
3. Click **New Role**.
4. In Name, provide a name for the role, and then click **Apply**.
5. Click **Members**, and then click **Add Roles**.
6. Select the roles to include and click **Add**.
7. Click **OK**.
8. As shown next, select **Inherit membership from children** and **Use intersection of role membership**, as the case may be.
9. Click **Save**.

Figure 8 Create an Administrator Defined Role

Filtering LDAP Users and Groups to Integrate

An LDAP directory can contain many users and groups. In many cases, you'll want to use only a subset of the users and groups. When using TIBCO Domain Utility to create an administration domain that is integrated with an LDAP directory server, you typically specify a search filter to retrieve only the LDAP users and groups that are relevant for the applications managed in TIBCO Administrator.

Even if you choose not to automatically create a role for each LDAP group, you'll want to use a search filter to return only limited LDAP groups. The synchronization feature only acts against the LDAP groups returned by this filter.

User and group search filters are written using the syntax defined in RFC 2254 *The String Representation of LDAP Search Filters*. See the *TIBCO Runtime Agent Domain Utility User's Guide* for more information and examples.

TIBCO Administrator allows you to modify the choice of groups (and optionally their descendants) to synchronize. After you save the customizations, they are used for future syncs, rather than those set when the administration domain was created.

Selecting LDAP Groups to Synchronize in TIBCO Administrator

In addition to filtering the groups in Search filters specified in TIBCO Domain Utility, you can further limit the LDAP groups that need to be synchronized.

At any time you can force the administration server to immediately synchronize with the LDAP server. You can:

- Specify that all groups be synchronized
- Specify that selected groups be synchronized. You can add or remove groups from the selection list and specify, for each group, whether the synchronization should include subgroups (descendants) or not.



When you change the list of groups to synchronize, the change becomes the default list and is used the next time an automatic synchronization is scheduled.


- Use the `CorpRoleSynchronizer` command line utility to synchronize your domain with all LDAP groups in the LDAP directory associated with the domain.

To Change Synchronization Criteria

After you select and save a subset of groups to synchronize, only those groups are kept in sync with changes in the LDAP directory. Also, the changed synchronization criteria is used the next time an automatic synchronization is triggered.


1. In the TIBCO Administrator GUI, in the left panel under User Management, select **Roles**.
2. Click **Select LDAP Groups**.
3. Click **Selected Groups**, and then click **Add**.

Figure 9 Select Groups to Synchronize



4. Select the groups to synchronize and click **Add**. For example, the next diagram shows that the TEAK group has been added.

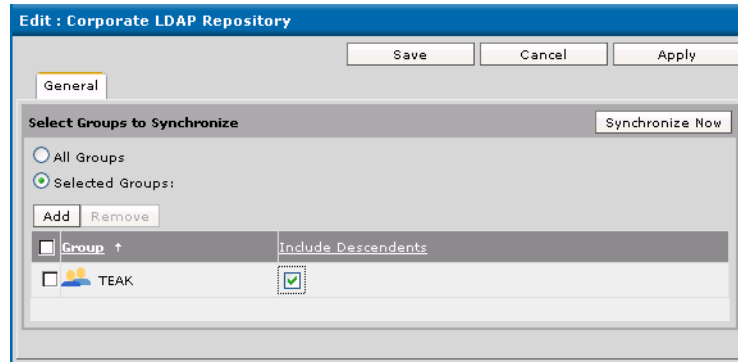
Figure 10 Add Group



5. Click **OK**.

- To include sub groups, select the **Include Descendents** check box next to each group. For example, the next diagram shows the selection. All groups in the list and subgroups will be synchronized.

Figure 11 Include Descendents Check Box



- You are now ready to synchronize. There are two choices:

If you click **Synchronize Now**, TIBCO Administrator blocks while the synchronization operation is performed. When control returns, you must click **Save** to reuse the settings.



If you click **Cancel**, the settings are lost and the next automatic synchronization will use the previously saved settings.

If you click **Save**, the synchronization operation occurs in the background. That is you can access other TIBCO Administrator screens while the synchronization operation is performed. You must refresh your Browser to see results. The automatic synchronization operation will use these settings the next time it is invoked.

To do an Immediate Synchronization

When you do an immediate synchronization, the criteria for automatic synchronizations are not changed. An immediate synchronization uses the criteria from the previous synchronization operation.

- In the TIBCO Administrator GUI, in the left panel under User Management, select **Roles**.
- Click **Select LDAP Groups**.
- Click **Synchronize Now**.
- Click **Save**.

CorpRoleSynchronizer Command Line Utility

The CorpRoleSynchronizer command line utility syncs an administration domain with its associated LDAP directory. The sync occurs based on the search criteria for LDAP groups that was defined when the administration domain was created.

Your administration domain may be out of sync because the auto sync settings for the domain were not enabled when using TIBCO Domain Utility to configure the domain, or because significant changes have been made to the LDAP directory since the last automatic sync and you do want to wait for the next auto sync cycle to occur, and you do not want to do an manual sync from the TIBCO Administrator GUI.

The utility is located in the *TIBCO_HOME/tra/<version>/bin* directory. The domain name you provide must have been configured to use an LDAP directory server.

```
C:\tibco\tra\<version>\bin>CorpRoleSynchronizer -h
USAGE: CorpRoleSynchronizer -domain <domain> [-h|-?]
where
    -domain <domain> - Name of a domain (case sensitive)
    -h or -?         - prints this help information
```

Note that the utility must be started with an option or an exception will result.

A summary of results is provided in the console where you launched the utility and in the TIBCO Administrator log file. Note that the command must be started with an option or an exception will result.

Setting the Maximum LDAP Objects to Return After a Search

By default, the maximum number of LDAP objects returned to TIBCO Administrator for a search is 10000. You can override the default by adding or changing the **DomainUsersSearchLimit** property in the *AuthorizationDomain.properties* file. The file is located in the *TIBCO_HOME/tra/domain/domain* directory.

Note that the client-side search limit is overridden by the LDAP server search limit. You may also have to change the corresponding setting on the LDAP server.

Pre Loading User Objects

When an application is running, a new user profile is created when a user first accesses the application. At that time, the default objects for the user's profile are created in the administration domain's LDAP repository.

Because of this activity, the administration server can become overloaded if many new users access the application in a short period of time. For example, this is especially applicable to TIBCO PortalBuilder applications directly after they go into production. In these situations, performance will be improved if you run the `CorpUserSynchronizer` utility.

Pre loading users is generally done before an application goes live, but can be done at any time when a mass import of users would be useful.

CorpUserSynchronizer Command Line Utility

The utility is located in the `TIBCO_HOME/tra/5.9/bin` directory. The domain name you provide must be configured to use an LDAP directory server.

```
C:\tibco\tra\5.9\bin>CorpUserSynchronizer -h
USAGE: CorpUserSynchronizer -domain <domain> [-clean] [-h|-?]
where
    -domain <domain> - Name of a domain (case sensitive)
    -clean           - this will remove the users from domain that
are no longer present in Corp LDAP
    -h or -?        - prints this help information
```

Note that the utility must be started with an option or an exception will result.

Importing Large User Bases

If you have a large number of users to import, the import operation may fail after processing only part of the LDAP directory's contents. This happens because of low default values in certain LDAP directory server settings.



It is assumed that your LDAP directory servers (primary and referral servers) are appropriately configured so that LDAP search queries return all matching users.

- For iPlanet and Sun ONE Directory Servers, the **timelimit**, **sizelimit** and **look-throughlimit** are appropriately set.
- For Microsoft Active Directory, the `NTDSUTIL` utility is used to configure the appropriate settings. See the Active Directory documentation for instructions.

A summary of results is provided in the console where you launched the utility and in the TIBCO Administrator log file.

LDAP Synchronization Optimistic Option

For both automatic and manual synchronization, an *optimistic* option is available to improve performance. This parameter is valid for LDAP directory servers with or without referrals.

The membership list of an LDAP group contains users, other groups (called sub groups) and potentially other LDAP entry types such as computers or networks. The group membership list does not distinguish members and only stores the DN for each of these entries. While computing the entire group membership, TIBCO applications and services must distinguish among the entries in the membership list. The sub group entries in the list are identified by their DNs, but for other entries the application must individually retrieve the LDAP entry for each DN. This is potentially a performance issue.

If the optimistic option is used, TIBCO applications and services assume that all entries that are not sub groups are valid users, and does not query the LDAP directory server. Using the optimistic option reduces the number of queries, and thus improves performance.

To Set the Optimistic Option

The optimistic option is set when creating an administration domain that integrates with an LDAP directory server. See the *TIBCO Runtime Agent Domain Utility User's Guide* for details.

Prerequisites for Using the Optimistic Option

Before using this option, you must ensure that the following are true (otherwise unpredictable results may occur):

- The static membership list of any group in the LDAP directory server contains only sub-groups and valid (existent) users. If it contains other types of items such as computers or networks, the algorithm will work only if there are no users in the LDAP directory server that have the same name as any of these other items.
- The static membership lists contain only users that are integrated based on user search filters specified in the LDAP setting for the administration domain. The static membership lists do not contain users that are unreachable because of LDAP referrals for which credentials are either not provided or are provided incorrectly.
- The static membership lists do not contain members that are aliases of unreachable users or sub groups.

Chapter 4 **Managing Users and Roles**

This chapter explains how to create and configure users and roles in standard administration domains. If you have integrated TIBCO Administrator with an LDAP directory server, see [Chapter 3 on page 36](#).

Topics

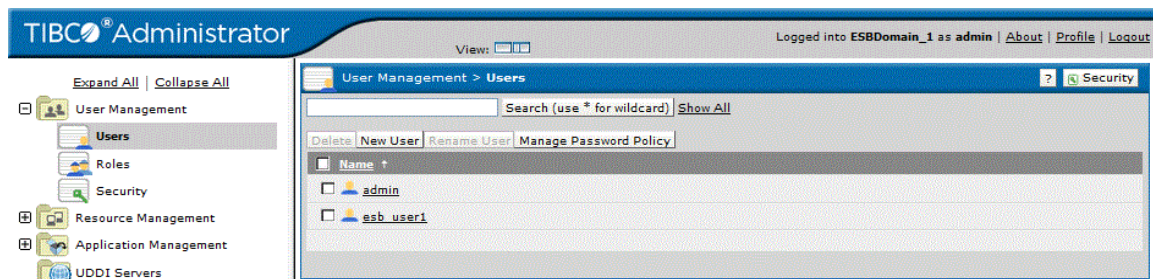
- [User Management Overview, page 58](#)
- [Managing Access Rights, page 60](#)
- [Adding Users, page 64](#)
- [Assigning Role Membership to Users, page 65](#)
- [Removing Role Membership for a User, page 67](#)
- [Assigning Permissions to Users, page 68](#)
- [Changing or Resetting Passwords, page 69](#)
- [Deleting Users, page 71](#)
- [Renaming Users, page 72](#)
- [Changing Domain Administrator User Credentials, page 73](#)
- [Managing the Password Policy for an Administration Domain, page 75](#)
- [Creating a Role, page 81](#)
- [Adding or Removing a User from a Role, page 82](#)
- [Removing a Child Role from a Parent Role, page 83](#)
- [Assigning Permissions to Roles, page 84](#)
- [Users Dialog, page 86](#)
- [New User Dialog, page 87](#)
- [Roles Dialog, page 89](#)
- [New Role Dialog, page 90](#)
- [Profile Dialog, page 93](#)

User Management Overview

The User Management module allows you to manage users, roles and security for an administration domain. You specify users and roles using the Users and Roles consoles and give users and roles access to individual components in the TIBCO administration domain using the Security console. The console is explained in [Chapter 5 on page 94](#).

- The Users console allows privileged users to create users. It also allows users who have the Administer permission to grant other users permission to read, write, or administer any of the consoles or other resources in TIBCO applications and repositories.
- The Roles console allows privileged users to create roles. The purpose of roles is to allow an administrator to group users, and then simultaneously give (or withdraw) permission for that group of users to use consoles, applications, or application repositories.
- Command line tools can also be used to create and update users and roles. See `ImportDomainSecurity` and `ExportDomainSecurity` in the *TIBCO Administrator Server Configuration Guide*.

Figure 12 User Management



Using the Guest Role

The TIBCO Administrator GUI role console displays a **Guest** role as a sibling to authenticated users. Privileged users use this role to designate certain consoles or resources in TIBCO applications that do not require log-on.

Searching for Users

Most TIBCO Administrator GUI consoles include a search function. The search function works the same for each console. For example, for the Users console:

- Typing a user's name and clicking Search displays only that user.
- Using the wildcard character (*) allows you to indicate zero or more characters. For example, if users Ravi, Rachel, and Tara are defined, *ra* returns Ravi, Rachel, and Tara, ra* returns Ravi and Rachel, and ra does not return anything.
- Search is case sensitive with the exception of searching users.
- Typically, search can be based on data that appears in any of the columns.

Selecting Items

There are multiple ways to select items in the TIBCO Administrator GUI. For example, for the Users console:

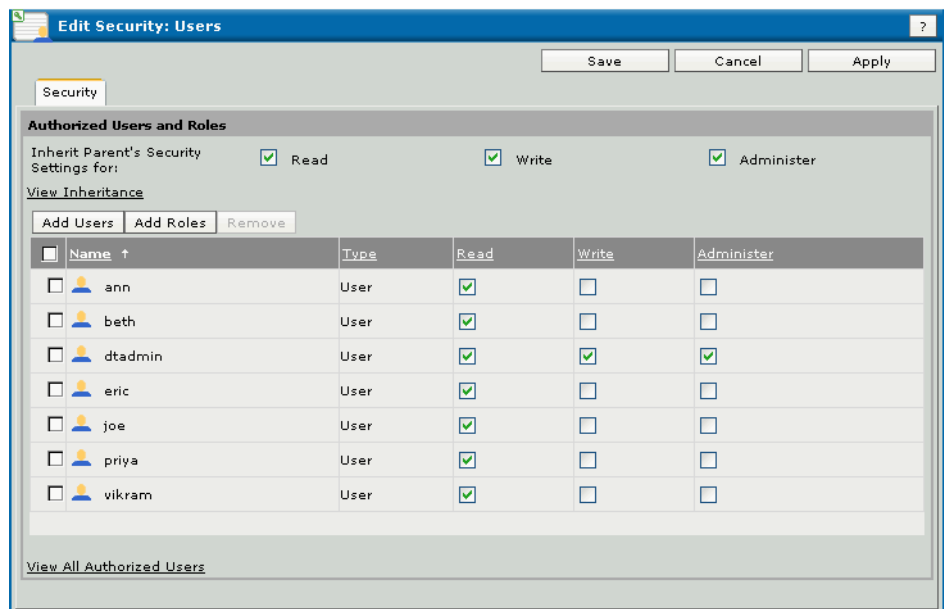
- Clicking Name displays users in ascending or descending alphabetical order.
- Clicking the check box next to Name selects or clears all users.
- Shift-clicking selects check boxes for multiple sequential users.
- Clicking selects check boxes for multiple non sequential users.
- Clicking a user's name (not a check box) displays a dialog where that user's properties can be viewed or edited.
- Sorting a column clears the selection.
- A selection can span pages (in cases where the table or list is paged).

Managing Access Rights

The TIBCO Administrator GUI provides Read, Write and Administer permissions for items that display in the Security console tree. The domain administrator user, a Super User or a user with Administer permission sets permissions. See [Chapter 5 on page 94](#) for security details.

Access permissions are set using the Security console that is available for each TIBCO Administrator GUI console. For example, the following diagram shows the Security dialog for the Users console. All users have Read access and the dtadmin user has Write and Administer access. Read access allows users to view the Users console. Write access allows users to change domain data through console and Administer access allows users to set access permissions for other users. Access rights are inherited down the security console tree unless specifically overridden.

Figure 13 Edit Security



Read Access

A user with read access to a resource can view that resource.

Read access provided at the top level folder in the Security console tree allows you to:

- view all users and roles defined in the domain.
- view software installed in the domain. You can export inventory information to a file.
- view the machines that are part of the domain. You can see details about the machine, including alerts, running processes and rulebases and events defined for each machine.
- view the applications loaded into the domain. For process engines or service instances, you can determine the machines on which each is running, the rulebases and events defined for each and properties set for each.
- view the state (started or stopped) of each service instance and process engine.
- query log files for each service instance and process engine.
- view plug-ins and their properties.

Read access does *not* allow you to:

- delete or otherwise modify items in any of the consoles.
- start or stop applications.
- view the Security console or permissions set on items.

Write Access

A user with write access to a resource can modify that resource. Write access to a resource implies read access to that resource.

Write access is typically assigned to developers who load, deploy and monitor applications. Developers can be assigned write access to just the applications and application repositories they are responsible for or can be given write access to all applications. Developers typically need not have access to the User Management module.

Write access provided at the top level folder in the Security console tree allows you to:

- create or delete users.
- create or delete roles. You can remove a role from its parent. You can also add users and roles to other roles.
- manage installed software. This includes removing, enabling or disabling software, and adding custom software.
- configure monitoring for machines in the domain by adding rulebases, events, or both.

- manage applications including adding, removing, deploying, and upgrading them.
- manage services and processes by changing global variables.
- create monitoring rulebases and events for service instances and process engines.
- start and stop service instances and process engines.
- query log files for each service instance and process engine.
- use installed plug-ins.
- write data to the domain and application repositories.

Write access does *not* allow you to:

- assign permissions to users or roles.
- add a machine to a domain.
- reset another user's password.
- view the Security console or permissions set on items.

Administer Access

A user with administrator access to a resource can assign permission to other users and roles to access that resource.

When permissions are set on a folder, they automatically apply to all items in the folder. This behavior can be changed such that Administer permission is only granted to some items in a folder. See [Chapter 5 on page 94](#) for details.

- The Administer permission implies Read permission but does not imply Write permission.
- A Super User or the domain administrator user can assign Administer permissions to others.

Super User Access

The domain administrator user can assign super user access to other users by adding them to the list of super users. A Super User has Read, Write and Administer permissions to all resources in the administration domain without explicitly having been granted those permissions. This allows the user to:

- manage all parts of domain.
- add a machine to a domain.
- reset another user's password.

A Super User can also add other users to the list of super users.

Adding Users

This section explains how to create users in an administration domain that is not integrated with an LDAP server. Each user is automatically a member of the Authenticated Users role.

After you save, the new user is displayed in the Users console. A user name is always created in lower case. This means you *cannot* have users named mike and Mike. All spaces before and after the name are removed. Spaces inside the name are retained.

To Add a User

1. Select **User Management > Users** in the left panel.
2. Click **New User**.
3. Supply the user name and click **OK**.
4. Click **set** to provide a password for the user. After entering the password, click **OK**.
5. Click **Save**.

See Also

New users can have permissions, roles or both assigned before saving. See [Assigning Role Membership to Users on page 65](#).

Privileged users can assign permissions to another user. See [Assigning Permissions to Users on page 68](#).

Assigning Role Membership to Users

The Add Roles dialog displays when you assign role membership to users or grant Security settings to a user or role.



Before you can assign role membership to users, you must have created the roles. See [Creating a Role on page 81](#).

In most cases, you assign users to a role when creating or editing the role. However, when you create a new user, or when several roles change for a single user, it makes sense to assign role membership to a user.

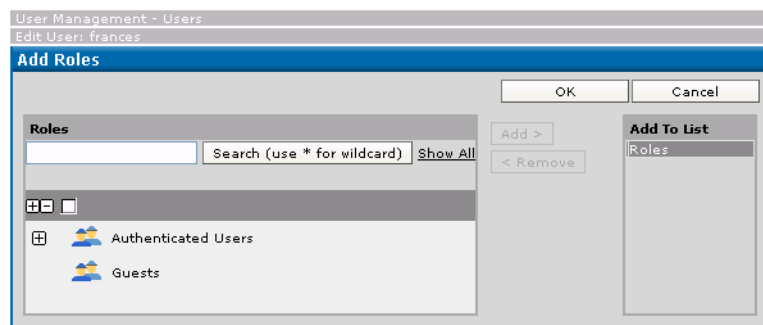
To Assign Role Membership to Users

1. Select **User Management > Users** in the left panel.
2. Select a user by clicking the name in the right panel.
3. Click **Role Membership**.
4. Click **Add Roles**.
5. Click one or multiple role names.
6. Click **OK**.
7. Click **Save**.

Add Roles Dialog

The Add Roles dialog displays when adding a user to a role.

Figure 14 Add Roles Dialog



The Add Roles window supports two different displays:

- In the tree hierarchy (the default display), you can expand individual roles using the + next to them or expand the entire tree using the + in the title bar.
- If you click **Search**, a flat display of all roles is presented. In this display (which potentially requires a lot more screen real estate), you can Shift-click to select sequential roles.

Add role membership for one or more roles using one of these options:

- Add membership to a single role for the user by clicking the role name, and then click **Add**.
- Select multiple roles by clicking the check boxes, and then click **Add**.

Removing Role Membership for a User

In most cases, you remove users from a role when editing the role. However, when several roles change for a single user, it makes sense to edit the user and remove role membership explicitly.

If an administration domain is integrated with an LDAP directory server, some groups may be organized as group-synchronized roles. You cannot add or remove users from group-synchronized roles.

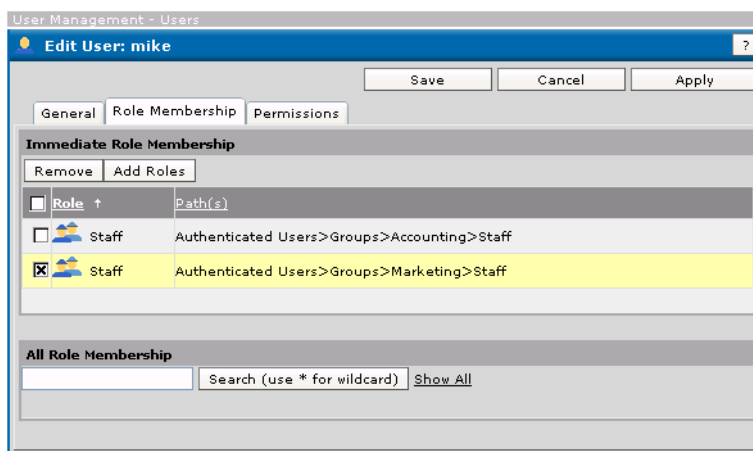
To Remove Role Membership for a User

1. Select **User Management > Users** in the left panel.
2. Select a user by clicking the name in the right panel.
3. Click **Role Membership**.
4. Select the checkbox next to the role(s) from which this user will be removed.
5. Click **Remove**.
6. Click **Save**.

Edit Roles Dialog

The edit roles dialog allows you to remove a selected user from a role.

Figure 15 Edit Role Dialog



Assigning Permissions to Users

Users that have Administer access can assign other users permissions to TIBCO Administrator GUI consoles, applications or application repositories when they create the user or at a later time.

To Assign Permissions to Users

1. Select **User Management > Users** in the left panel.
2. Select a user by clicking the name in the right panel.
3. Click **Permissions**.
4. Under Authorization for, expand the TIBCO Administrator and Data Access folders and assign permissions.
5. Click **Save**.

See Also

See [Managing Access Rights on page 60](#) for information about permissions.

See [Chapter 5, Granting Security Access to Objects, on page 94](#) for more information.

Changing or Resetting Passwords

A user must supply a current password to change it.

Only a member of the Super User role can reset another user's password. This typically occurs if a user forgets his or her password.



If you integrate with an LDAP directory server to manage users, you cannot change passwords for LDAP users in the TIBCO Administrator GUI. Contact your LDAP administrator if you need to change passwords for LDAP users.

To Change Your Password

1. Select **Profile** in the top, right panel.
2. Enter your current password and click **OK**.
3. Click **Change** next to the password field.
4. Enter the new password twice.
5. Click **OK**, and then click **Save**.
6. Redeploy all applications with the following characteristics:
 - the application is deployed with http, https, or rvtransport options for application data.
 - the user whose password you just changed is used to access the repository server for application data.

Before you redeploy, you must manually supply the new password in the application's configuration panel in the **Advanced** tab of the Edit Application Configuration dialog for that application. See [TIBCO BusinessWorks and Adapters Deployment Repository Instance on page 201](#) for more information.

You also can use the command line tool `redeployAllAppsForUsers` to redeploy all the applications at once. Refer to *TIBCO Runtime User's Guide* for detail.

To Reset a User's Password

You must be a Super User to reset a user's password.

1. Click **User Management**.
2. Click **Users**.
3. Click the user name.

4. Click **Change** next to the password field.
5. Enter the new password twice.
6. Click **OK**, and then click **Save**.



The user whose password was just reset must redeploy all applications with the following characteristics:

- the application is deployed with http, https, rv, or ems transport options for application data.
- the user whose password you just changed is used to access the repository server for application data.

Before you redeploy, you must manually supply the new password in the application's configuration panel in the **Advanced** tab of the Edit Application Configuration dialog for that application. See [TIBCO BusinessWorks and Adapters Deployment Repository Instance on page 201](#) for more information.

See Also

See [Changing Domain Administrator User Credentials on page 73](#).

Deleting Users

A user with Administer permissions to the User console can delete other users from the Users console. If only one user with Super User privileges exists, you will be blocked from deleting that user. It is therefore only possible to delete the domain administrator user, if you have created additional Super Users.

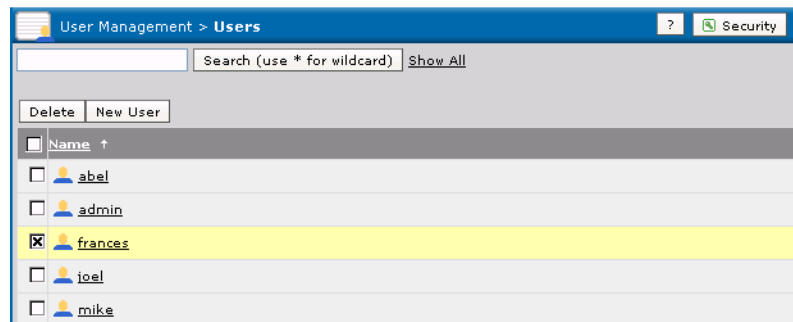


If you integrate with an LDAP directory server to manage users, you cannot delete LDAP users from this console but must delete LDAP users directly in the LDAP console. Contact your LDAP administrator for assistance.

To Delete a User

1. Choose **User Management > Users** in the left panel.
2. Click the selection box next to one or more users.
3. Click **Delete**.

Figure 16 Delete a User



4. Click **OK** when prompted to confirm the delete operation.

Renaming Users

A Super User or user with Administer permissions to the User console can rename users. When you rename a user, the new user has the same privileges as the old user.



If you integrate with an LDAP directory server to manage users, you cannot rename LDAP users in the TIBCO Administrator GUI. You need to rename the LDAP users in the LDAP console. Note that when the system synchronizes after a rename, the privileges assigned to the user (role assignment and access rights) are lost.

To Rename a User

1. Choose **User Management > Users** in the left panel.
2. Select the check box next to the user name to rename.
3. Click **Rename User**.
4. Change the name in the editable User Name field.
5. Click **Save**.
6. Redeploy all applications with the following characteristics:
 - the application is deployed with http, https, or rv transport options for application data.
 - the user that you renamed is used to access the repository server for application data.

Before you redeploy, you must manually supply the new username in the application's configuration panel in the Advanced tab of the Edit Application Configuration dialog for that application. See [TIBCO BusinessWorks and Adapters Deployment Repository Instance on page 201](#) for more information.

Changing Domain Administrator User Credentials

The domain administrator user is the original user created when the administration domain was created. You can:

- Change the domain administrator user's password. You must first change the password in the Users console, and then use TIBCO Domain Utility to change the domain credentials. You must use Domain Utility to change domain administrator credentials on each machine in the domain by running the Domain Utility on each machine.
- If you want another user to be the domain administrator user, that user must be defined in the Users console and be a member of the Super User role. If your domain is integrated with an LDAP directory, the user must be defined in the LDAP directory and be a member of the Super User role. After verifying this, use TIBCO Domain Utility to change the domain administrator credentials.



The administration server must be running when using TIBCO Domain Utility to change the domain administrator credentials. You could lose the administration rights for your domain if the server is stopped while changing domain administrator credentials.

To Change Domain Administrator User Credentials

1. Choose **User Management > Users** in the left panel.
2. Click the domain administrator user name.
3. Click **change** and set the new password.
4. Click **OK**.
5. Click **Save**.
6. On *each* server and client machine for the administration domain, start TIBCO Domain Utility and do the following:
 - a. Select **Server Settings > Change Domain Credentials** and click **Next**.
 - b. Select a domain.
 - c. In the screen that displays, provide values in each field and click **Next**.
 - d. Click **Next** to apply the values.
 - e. **(Server machine only)** Restart the administration server.
 - f. Restart the TIBCO Hawk agent.

7. Log into TIBCO Administrator GUI again using the new domain administrator's credentials and redeploy all applications with the following characteristics:
 - the application is deployed with http, https, rv, or ems transport options for application data.
 - the domain administrator's credentials are used to access the repository server for application data.

Before you redeploy, you must manually supply the new username and password in the application's configuration panel in the **Advanced** tab of the Edit Application Configuration dialog for that application. See [TIBCO BusinessWorks and Adapters Deployment Repository Instance on page 201](#) for more information.

You also can use the command line tool `redeployAllAppsForUsers` to redeploy all the applications at once. Refer to TIBCO Runtime User's Guide for detail.

See Also

See Changing Domain Credentials in *TIBCO Runtime Agent Domain Utility User's Guide* for more information.



TIBCO Admin 5.9 supports DB2 9.7 HA database for new domains. For domains that are migrated, you need to recreate the domain using Admin 5.x

Managing the Password Policy for an Administration Domain

A password policy can be set for an administration domain when it is created. Users with Write permission to the Users console can change the password policy by editing the existing password policy, creating a new password policy or removing the password policy.

For an introduction to the password policy, see the *TIBCO Administrator Server Configuration Guide*.

When setting a password policy, you can either use one of the XML templates provided in the `TIBCO_HOME/tra/<version>/config/security` folder or create one of your own. After you load the password policy file, it is stored in the administration domain. The `PasswordPolicy.xsd` file and the following template files are available:

- `DefaultPolicy.xml`
- `NormalPolicy.xml`
- `StrongPolicy.xml`

Once created, the password policy applies to all users and groups in the administration domain. You should use an LDAP directory server if you wish to customize password policies for different users and groups.



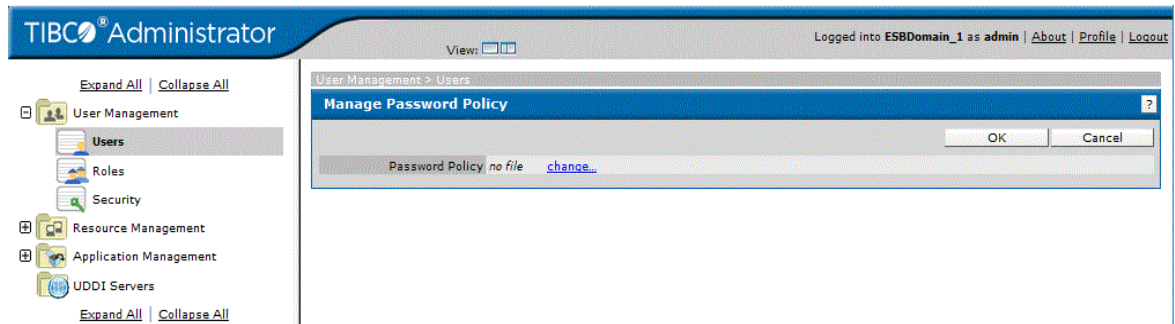
The password policy management operations are now recorded in the `audit.log` file. For more details, see *TIBCO Administrator Server Configuration Guide*.

Modifying the Password Policy

To modify the password policy for an administration domain:

1. Copy an XML template file in `TIBCO_HOME/tra/<version>/config/security` to another directory.
2. Using an XML editor, modify the XML password policy file for your environment. See [Password Policy Elements on page 76](#) for details.
3. Log in to the TIBCO Administrator GUI using a user account that has Write permissions to the Users console.
4. Expand **User Management** and select **Users**.
5. Click **Manage Password Policy**.

Figure 17 Manage Password Policy



6. Click **change** and in the dialog that appears, click **Browse**. Navigate to the XML password file you modified earlier. Click **OK**.
7. Click **OK**.

Removing the Password Policy

After you remove the password policy for a domain, no policy is enforced for passwords. Users can be created without assigning passwords.

To remove the password policy for an administration domain:

1. Log in to the TIBCO Administrator GUI using a user account that has Write permissions to the user management users console.
2. Expand **User Management** and select **Users**.
3. Click **Manage Password Policy**.
4. Click **remove**.
5. Click **OK**.
6. Click **OK**.

Password Policy Elements

The password policy elements found in the XML template files provided in the `TIBCO_HOME/tra/<version>/config/security` folder are explained next.

Table 3 Password Policy Elements

Element	Description
PolicyName	Name assigned to the policy.

Table 3 Password Policy Elements

Element		Description
SaveHashMode		<p>True or false. If true (the default), the password will be hashed using the SHA1 algorithm. If the password is hashed, it cannot be decrypted.</p> <p>If false, the password will be encrypted using 3DES-CBC with a 192-bit key. If the password must be decrypted at runtime, the password cannot be hashed.</p>
PasswordLength		
Min		Required. The minimum number of characters a password can have.
Max		Optional. The maximum number of characters a password can have.

Table 3 Password Policy Elements

Element	Description
PasswordComplexity	
MinRequirement	Optional. The number of Requirement elements (Numbers, SpecialCharaters, UpperCase, LowerCase) that must be used in the password.
ContainCurrentPass word	Optional. True or false. If true, a user's existing password cannot be contained in a new password when changing the password. If false, an existing password can be included in the newly changed password.
ContainUserId	Optional. True or false. If true, the user's account name cannot be part of the password. If false, the account name can be part of the password.
NoSpace	Optional. True or false. If true, spaces are not allowed in the password. If false, spaces are allowed in the password.

Table 3 Password Policy Elements

Element	Description
Requirement	
Numbers	Optional. If included, using at least one number between 0 and 9 inclusive will meet the requirement in the password.
SpecialCharacters	Optional. If included, using at least one of the following special characters will meet the requirement in the password: , ! @ # \$ % ^ & * () _ + ~ - = \ ` { } [] : " ; ' < > ? , . /
UpperCase	Optional. If included, using at least one character in upper case will meet the requirement in the password.
LowerCase	Optional. If included, using at least one character in lower case will meet the requirement in the password.
PasswordAging	
Min	Optional. The minimum number of days a password must be kept before the user can change it.
Max	Required. The maximum number of days a password can be kept before the user must change it.

Table 3 Password Policy Elements

Element	Description
ForceInitialPasswordChange	Optional. True or false. If true, the user must change the password upon first login and when the password is reset. If false, the user need not change the password when initially logging in or when the password is reset.
AccountLockout	Optional. The number of failed login attempts a user can make before his or her account is disabled.
PasswordHistory	Optional. The number of unique new passwords that must be used before an old password can be reused.

Creating a Role

A role can be created directly under the Authenticated Users role or under another role (as a child role). Adding child roles creates role hierarchies that can be used to organize roles for searching purposes.

Role names are case sensitive. This allows you to create roles named marketing and Marketing.



You can rename Authenticated Users if desired (for example, to XYZ Incorporated Users), but you cannot delete the role.

To Create a Role

1. In the left panel, choose **User Management > Roles**.
2. Expand the hierarchy and select the check box to the left of the role that you want to use as the parent role for the new role (either Authenticated Users or another role previously created).
Do not click the role itself for this task.
3. Click **New Role**.
4. Specify the role name.
5. Click **Apply** to use the other tabs.
6. Click **Save** to create the role.

See Also

You can add members to the role while creating it, or at a later time if you do not currently want to add members. See [Adding or Removing a User from a Role on page 82](#) for more information.



If you integrate with an LDAP directory server, each group in the LDAP directory becomes a group-synchronized role. See [Chapter 3, Integrating TIBCO Administrator with an LDAP Directory Server, on page 36](#) for more information.

Adding or Removing a User from a Role

There are two situations under which you may be prompted while modifying the members of a role:

- When you select the Security console and click **Super Users**, you can add (or remove) users to that role.
- When you select the Roles console, you can add or remove users by clicking the selected role. You can also add other roles in the membership list of a role. You cannot remove users (or roles) from the membership list of the **Authenticated Users** role. This is because all valid users in an administration domain are implicitly considered members of this role.

Edit Role > Members display shows all immediate members of that role. Immediate members are the direct members of the given role. Members that are based on LDAP groups or inherited members through the role hierarchy are not immediate members. To see all members, click **View All Members**.

To Add a User to a Role

1. In the left panel, choose **User Management > Roles**.
2. Click a role name.
3. Click the **Members** tab.
4. Click the **Add Users** button.

In the Add Users window that is displayed, click an individual user's name in the left panel to make that user a member, or select multiple users by selecting the check box.

5. Click **Add**.
6. Click **OK**.
7. Click **Save**.

To Remove a User from a Role

1. In the left panel, choose **User Management > Roles**.
2. Click a role.
3. Click the **Members** tab.
4. Select the user to remove and click **Remove**.
5. Click **Save**.

Removing a Child Role from a Parent Role

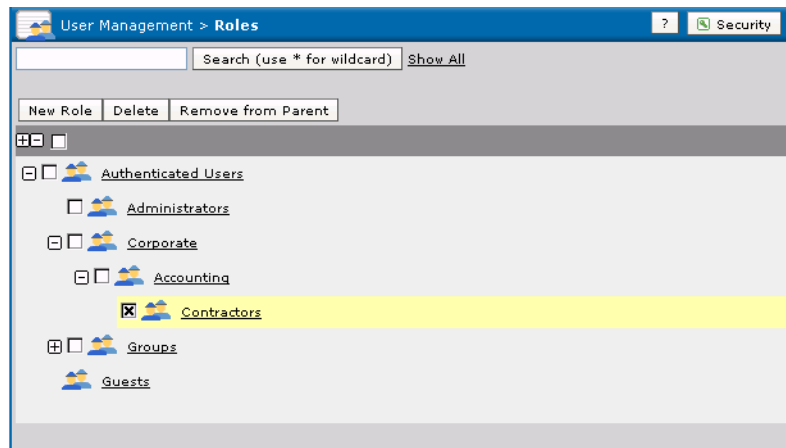
A child role is always created as a child of an existing role (which could be the Authenticated Users role or another role). Remove from Parent removes the selected role from its parent role. If the role has no other parents, it is made a child of the Authenticated Users role.

You can also perform this operation under the Roles Tree tab in the Edit Roles dialog by selecting the role(s) to remove and clicking **Remove**.

To Remove a Role from its Parent Role

1. In the left panel, choose **User Management > Roles**.
2. Expand the hierarchy and select the check box for the child role to remove.
3. Select the **Roles Tree** tab.
4. Click **Remove from Parent**.
5. Click **OK** in the popup that appears.

Figure 18 Remove a Role from Its Parent Role



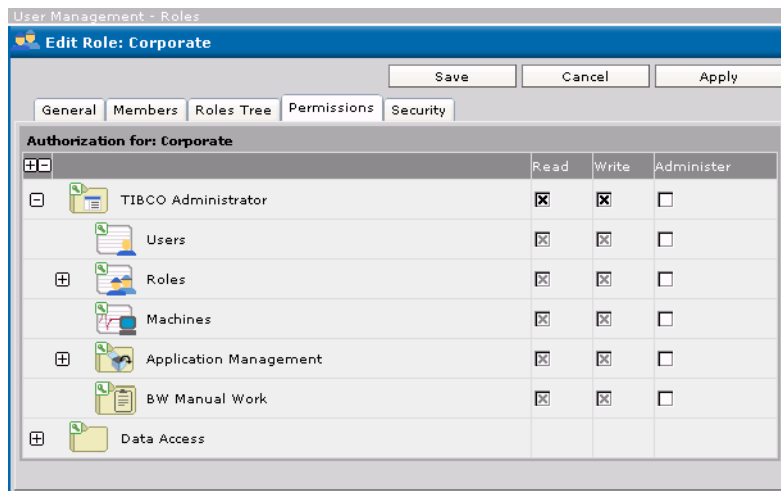
Assigning Permissions to Roles

Users that have Administer access can assign roles permissions to TIBCO Administrator consoles or Data Access when they create a role or at a later time.

The following diagram shows the Edit role dialog. All cleared check boxes allow you to give the role permissions directly.

All selected, gray boxes indicate that the role either inherits rights for those data from higher in the hierarchy or that the role is a member of a role for which access rights have been specified. Change authorization for the higher-level element or the role to change those access rights.

Figure 19 Assigning Permissions to Roles



To Assign Permissions to Roles

1. Choose **User Management > Roles** in the left panel.
2. Select a role by clicking the name in the right panel.
3. Click the **Permissions** tab.
4. Assign permissions.
5. Click **Save**.

See Also

See [Managing Access Rights on page 60](#) for a description of Read, Write and Administer permissions.

See [Granting Security Access to Objects on page 94](#).

Users Dialog

Security

Displays the Security dialog box where you can add, remove or view the users and roles that have permissions to use the module.

Search

To view a user whose name matches an exact search string, enter the search string, and then click the **Search** button. The user name will display.

You can use a wildcard character to indicate zero or more letters. For example, if users Ravi, Rachel, and Tara are defined, ***ra*** returns Ravi, Rachel, and Tara, **ra*** returns Ravi and Rachel, and **ra** returns nothing.

To return to a display of all users, click the **Show All** button.

Delete

Removes selected users. Note that you cannot recover a deleted user. See [Deleting Users on page 71](#) for details.

New User

Displays the New User dialog, which is explained in [New User Dialog on page 87](#).

Rename User

Allows you to rename a user. A user name must be selected to activate this button. See [Renaming Users on page 72](#) for more information.

Manage Password Policy

If you have Write permission, you can change the password policy that was set when the administration domain was created. See [Managing the Password Policy for an Administration Domain on page 75](#) for more information.

New User Dialog

The following tabs are available:

- [General Tab](#)
- [Role Membership Tab](#)
- [Permissions Tab](#)

General Tab

User Name

Provide a user name. Names are not case sensitive. That is, you cannot create users named mike and Mike.

Password

If you do not specify a password, a user can log in with a blank value in the password field.

If you are logged in as the domain administrator user or a Super User, the **change** link appears. Clicking the link displays a dialog where you can assign or change a password for a user.

If a user is created in an administration domain and the password policy is set as Strong Policy, after the administrator is restarted, it will prompt for changing the password when logging in with the newly created user for the first time.

Role Membership Tab

Immediate Role Membership

Immediate members are the direct members of the given role. Members that are inherited from child roles are not immediate members.

Add Roles

Displays the Add Roles dialog where you can select the role(s) to which this user is to become an immediate member.

Remove

Removes the user from the selected role(s).

All Role Membership

Displays the parent roles to which this user belongs, excluding the roles in which this user is an immediate member.

Permissions Tab**Super User**

Indicates whether the user is a Super User. Yes indicates the user is a Super User and No indicates the user is not. Click Security in the left column to access the Security console where Super Users can be defined.

Authorization for:

All white check boxes allow you to give the user access rights directly.

All grayed out boxes indicate that the user either inherits rights for those data from higher in the hierarchy or that the user is a member of a role for which access rights have been specified. If your account has Administer permissions, you can change these permissions directly or navigate to the higher-level element or the role to change those access rights.

See Also

See [Granting Security Access to Objects on page 94](#).

Roles Dialog

Search

Allows you to display only the roles that match a search criteria. You can use the * character as a wildcard.

New Role

Allows you to define a new role. You must select the role under which the new role will be created. See [New Role Dialog on page 90](#).

If TIBCO Administrator is integrated with an LDAP directory server, this button is not available.

Delete

Removes the selected role(s). You are prompted to confirm the removal.

Remove from Parent

Removes the selected role from its parent role. If the role has no other parents, it is made a child role of Authenticated Users.

Select LDAP Groups

This button is only available if your domain is integrated with an LDAP directory server. Allows you to select the LDAP groups to synchronize. See [Managing LDAP Users and Group-synchronized Roles on page 45](#) for details.

New Role Dialog

The following tabs are available:

- [General Tab](#)
- [Members Tab](#)
- [Roles Tree Tab](#)
- [Permissions Tab](#)
- [Security Tab](#)

General Tab

Name

Provide a name for the role. Roles names are case sensitive. That is you can define roles named Dev and dev.

Click Save if you are finished. Otherwise, click Apply and click another tab.

Description

Provide an optional description for the role.

Members Tab

Immediate Members

Immediate members are the direct members of the given role. These can be users or other roles. Members that are inherited from child roles are not immediate members.

Inherit membership from children

If selected, all members of this role's direct child roles are also members of this role. By default the check box is clear, meaning members of a child role are not automatically members in its parent role.

Use intersection of role membership

If selected, only members common to each of the roles specified in its immediate membership are members of this role.

Add Users

Displays the Add Users dialog where all users are listed. You can select the users that you want to be part of this role. Click the Add button after selecting users. Click OK to add the users to the role.

Add Roles

Displays the Add Roles dialog where all roles are listed. You can select the roles that you want to be child roles to this role. Click the Add button after selecting roles. Click OK to add the child roles to the role.

Remove

Removes the selected role from this role's immediate membership.

View All Members

Shows all members who are part of this role. This expands the membership of all roles in its immediate membership list and also inherits from child roles if required.

Roles Tree Tab**Parents**

Lists the parent roles for this role.

Children

Lists the child roles for this role.

Remove

Removes the selected child role from this role.

Add Roles

Displays a dialog box from which you can select roles to add as children of this role.

Permissions Tab

Authorization for

Allows you to set privileges for this role. This includes access to TIBCO Administrator panels, data access, and applications.

- **Read access** allows view access to all screens for this option, but does not allow users to delete or otherwise modify items in any of the screens. Read access does not include, for example, deleting users or starting or stopping applications.
- **Write** gives users full access to individual consoles. This includes deleting or modifying items, and starting or stopping applications. Write access implies read access.
- **Administer** allows users to assign privileges to other users. Administer access implies read access but does not imply write access.

Security Tab

For each role, you can authorize which users, roles, or both can access it and at what level, Read, Write or Administer. See [Security Dialog on page 103](#) for more information.

See Also

See [Managing Access Rights on page 60](#) for a description of Read, Write and Administer permissions.

See [Granting Security Access to Objects on page 94](#).

Note that the edit roles dialog is the same as the new role dialog.

Profile Dialog

The Profile dialog allows you to change your password. You must provide your existing password to access the profile dialog.

You can also set the Auto Refresh Interval. This sets the amount of time in seconds to elapse before the TIBCO Administrator GUI updates its consoles. The default is zero, which means updates display in consoles when they occur. The Auto Refresh Interval only pertains to screens that display the auto-refresh icon. The icon is shown in [Turning Auto Refresh On or Off on page 119](#).



A change to the Auto Refresh Interval takes effect only after you logout and login to the TIBCO Administrator GUI.

You can toggle auto refresh state on or off. See [Turning Auto Refresh On or Off on page 119](#).

Chapter 5 **Granting Security Access to Objects**

This chapter explains how privileged users can set security access for other users to TIBCO Administrator consoles, applications and repositories.

Topics

- [Security Overview, page 95](#)
- [Granting Super User Access, page 98](#)
- [Granting Access to an Object, page 99](#)
- [Managing Concurrent Access, page 101](#)
- [Edit Security Dialog, page 102](#)

Security Overview

The Security console allows privileged users to manage authorization for other users in the system. Security access determines whether a user can perform an operation on a specific resource in an administration domain.

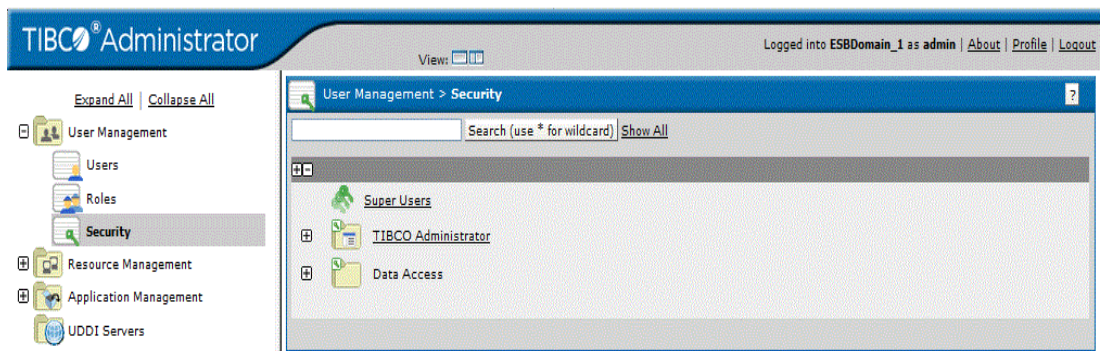


By default, the domain administrator user can manage security access and can assign other users the Administer permission to do the same. Users assigned to the Super User role, also have Administer access to the Security console. Other users cannot view the Security console.

Select **User Management > Security** to display the console. The console allows you to set security access for users and roles to TIBCO Administrator consoles and applications, and to repositories in the Data Access folder. Members of the Super Users roles are also assigned in this console. Each console or application has a small key icon associated with it. The green icon indicates the currently logged in user can assign others access permissions for the corresponding console.

See also, [Managing Access Rights on page 60](#) for an introduction to Read, Write and Administer permissions.

Figure 20 Security Overview



Security Console Tree

The Security console allows the domain administrator user, a member of the Super User role, or a user with Administer permission to assign access permissions to consoles and repositories. The Security console tree has two main folders, TIBCO Administrator and Data Access.

- The TIBCO Administrator folder displays the users, roles, and machines consoles. It also displays the Application Management console where permissions to each application loaded in the domain is set.
- The Data Access folder displays the repositories used by the domain and each application in the domain which use the server to store its application data, where access permissions can be set.

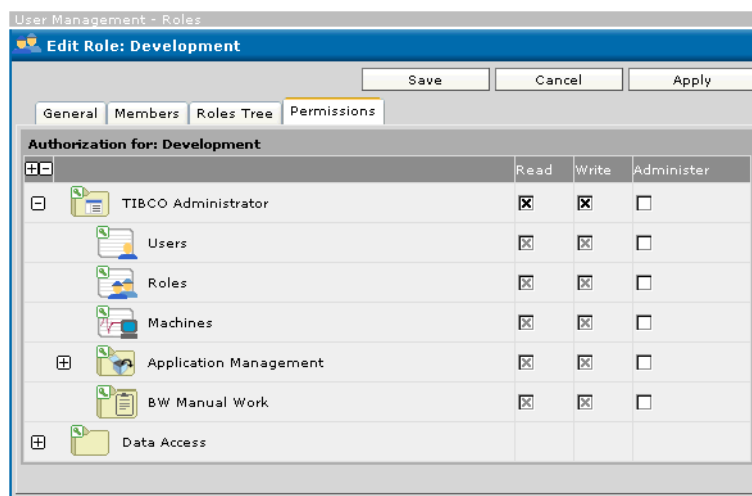
A Super User role member has read, write, and administer access to all TIBCO Administrator consoles and applications and data. A Super User role member can add other users to the Super Users role.

TIBCO Administrator Folder

When you specify permissions on a folder in the TIBCO Administrator folder, the permissions cascade down to all lower level items contained in that folder.

For example, the next diagram shows Read and Write permissions are assigned to the TIBCO Administrator module for the Development role. An **x** displays in bold for TIBCO Administrator and in grey for all objects contained in TIBCO Administrator. A bold x indicates permissions were assigned directly. Each grey x indicates the permissions were inherited, not assigned directly, or that the user is a member of a role for which permissions have been specified.

Figure 21 TIBCO Administrator Folder



You can set permissions for just one object by clearing permissions on the module or console that contains it. For the above example, to set Write access only to the Machines console, you must first clear the Write check box for TIBCO Administrator and then select the Write check box for the Machines console.

You can also break cascading permissions by changing the Inherit Parent's Security Settings option for an object. See [Inherit Parent's Security Settings on page 103](#) for details.

Data Access Folder

The Data Access folder contains folders that represent the domain repositories and application repositories. Each repository is displayed as a folder. Consequently permissions set on a repository apply only to that repository.

The Data Access folder allows privileged users to specify security access to its contained objects for other users. In particular, users who run the `appManage` command line utility and other command line utilities (see [Utilities, Plug-ins and Modules on page 23](#)) need Write access to these repository files.

A domain repository is used directly by the administration server, TIBCO Domain Utility and TIBCO Hawk agent. The domain repository contains data about the machines, registered software, users, roles, access control lists, application configurations and deployment history. In the case of a file-based domain, domain data is stored in the `SYS_domain.dat` and `AUTH_domain.dat` files.

Data stored in the `SYS_domain.dat` file is referred to as the administration domain while data stored in the `AUTH_domain.dat` file comprises the authorization domain. The authorization domain contains users, roles and data access ACLs. Everything else is stored in the administration domain: installed software, machines, applications, plugins, TIBCO Administrator ACLs, and so on. As such, the administration domain file is usually much larger than the authorization domain file.

TIBCO Administrator creates an application repository each time you deploy an application. An application repository contains information about the application's configuration and its deployment configuration using Rendezvous, http, or https as transport.



Do *not* use a text editor to change these repository files! You can potentially lose all domain information or deployment information for all applications.

Granting Super User Access

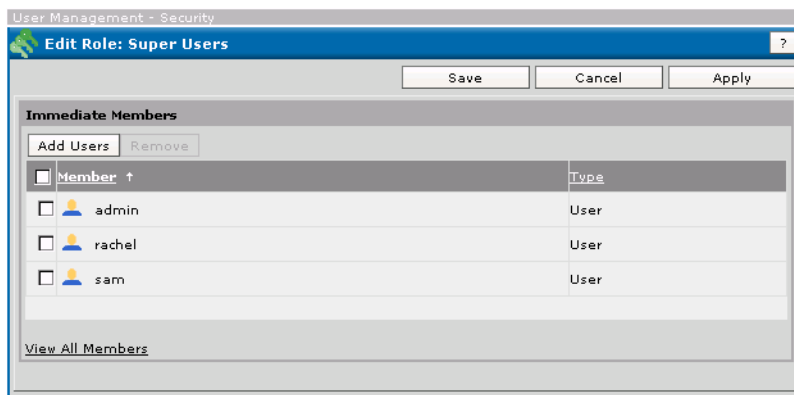
By default, the domain administrator user has Super User privileges. The domain administrator user can assign other users to the Super User role. Members of the Super Users role can assign other users to the role.

A user must exist before being added to the Super Users role.

To Grant Super User Access

1. Select **User Management > Security**.
2. Click **Super Users**.
3. Click **Add Users**.

Figure 22 Grant Super User Access



4. In the **Add Users** window, select the users you wish to add to the Super User role either by clicking on the underlined user name or by selecting one or multiple users, and then clicking **Add**.

To remove access, select the user and click **Remove**.

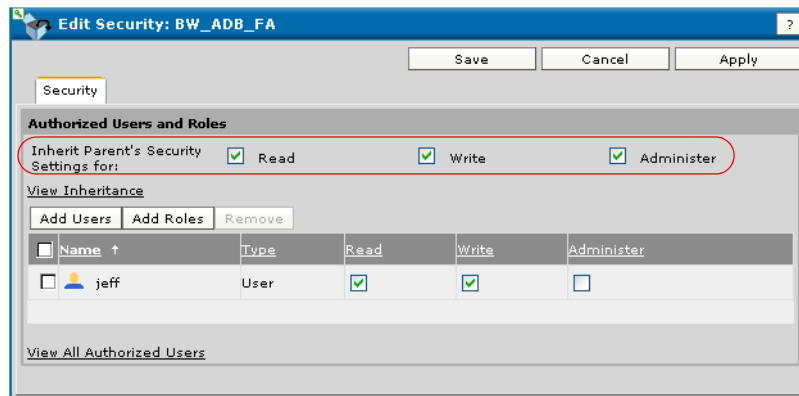
5. Click **Save**.

Granting Access to an Object

You must be the domain administrator user, a member of the Super User role, or have Administrator permissions to grant access to an object for other users and roles. Super User role members have access to all items in the Security console.

When you select a folder that contains other folders, the permissions you set for the top level folder apply to all contained folders. For example, if you are setting access to an application's folder, the permissions you set for a user or role apply to all of the contained folders. You can break this by changing which permissions an object inherits from its parent. The next diagram shows the option.

Figure 23 Granting Access to an Object



To Grant Access Permissions to an Object

1. Select **User Management > Security**.
2. Expand TIBCO Administrator and drill-down to select a folder or object name.
3. In the Security dialog, click **Add Users** or **Add Roles**.
4. Select users or roles and click **OK**.

To remove access, select the user or role and click **Remove**.

5. Select the permissions for the role or user.

If you don't want the object to inherit permissions from its parent object, clear the permissions listed under the Inherit Parent's Security Settings option.

6. Click **Save**.

See Also

See [Managing Access Rights on page 60](#) for a description of the Read, Write and Administer permissions.

TIBCO Administrator security access can be set when creating a user or role, or set on the item itself. See [Assigning Permissions to Users on page 68](#) and [Assigning Permissions to Roles on page 84](#).

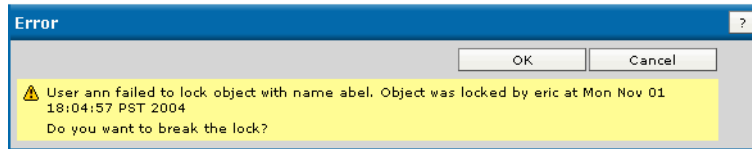
Managing Concurrent Access

Concurrent access means that two users are accessing a console in write mode at the same time.

TIBCO Administrator handles the situation as follows:

- When the second user accesses the console, that user is informed that another user has locked the console. The name of the user is provided. For example:

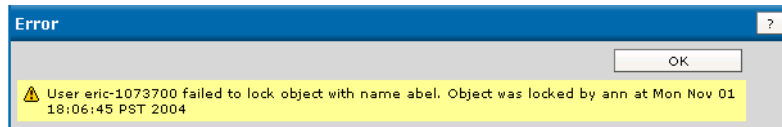
Figure 24 Object Locked



- The user is given the opportunity to break the lock. In general, you should break the lock only if you are sure that the first user is no longer working with that console and has, for example, exited uncleanly or gone on vacation.

If you break the lock and the other user is actually in the process of making changes to the same console, that user will then get a lock message when trying to save. That can potentially lead to confusing situations. For example, user eric would get the following message when attempting to save.

Figure 25 Message On Breaking Lock



Edit Security Dialog

This dialog lists the contents of the Security console, the TIBCO Administrator folder and the Data Access folders. Expand a folder and click an item to access the Security dialog where you can set access permissions to the item.

Security Dialog

Authorized Users and Roles

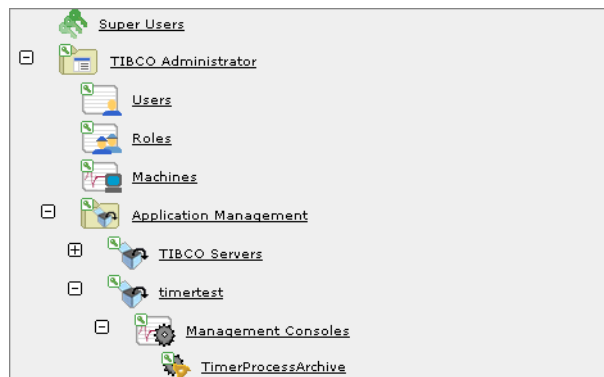
This pane provides a list of users and roles that can access the named resource in the top panel. For example, if you click **Machines**, a dialog displays where you can add the users or roles and set their permissions to access the **Machines** panel.

Inherit Parent's Security Settings

This option allows you to set whether the current item should inherit permissions from its parent in the Security tree.

For example, in the next diagram, a **TimerProcessArchive** application inherits all security settings from **Management Consoles**, which in turn inherits all settings from **timertest**, which inherits from **Application Management**.

Figure 26 *TimerProcessArchive Application Inheriting Security Settings*



By default, if you grant permissions to a user to the **Application Management** folder, that user also has the same permission to all items contained in the folder's tree. To give that user permission to only one application, or to only the applications, but not the **TIBCO Servers**, you must change the **Inherit Parent's Security Settings**.

You don't need to grant read permission at the top level since consoles which are parents to ones a user has read permission for will show up in the console tree. This is true even if the user does not have read permission to the parent. However, the parent will only display those children needed to navigate to consoles for which the user does have read access

View Inheritance

You must have Administer permissions for the object you are querying to view inheritance. The dialog that displays lists the name of the object that has permission, the access permission and the parent object that provides the permission setting for the object in question.

For example, the next diagram shows the permissions that are inherited by a service object. The Development role has Write permissions assigned on the service object's parent BW_ADB_FA. Three users also have permissions assigned to the service object at the top level TIBCO Administrator folder.

Figure 27 Authorization Inheritance

Name ↑	Read	Write	Administer	Authorized Object
Development	[x]	[x]	[]	TIBCO Administrator
abel	[]	[]	[x]	TIBCO Administrator
jeff	[x]	[x]	[]	TIBCO Administrator>Application Management>BW_ADB_FA
priya	[x]	[]	[]	TIBCO Administrator

Add Users

Displays the Add Users dialog where you can add users and then specify their access rights. Select one or more users from the Users list, and then click the Add button. Click OK to add the users to the new role. You can also change security settings. See [Permissions Tab on page 88](#) for a description of security settings.

Add Roles

Displays the Add Roles dialog where you can add roles and then specify their access rights.

Remove

Removes the selected user and role from the access list.

View all authorized users

Lists all users who have been authorized to use this console. This includes the users and roles membership inherited from parent resources, if any, and also those directly specified for the selected console. Security rights are listed for each user.



Role memberships include the users directly assigned to the role, the role memberships of any roles assigned to the role, and any role memberships inherited through the role hierarchy.

Chapter 6 **Managing Installed Software and Machines**

The TIBCO Administrator GUI Resource Management module allows you to create applications domains, and monitor and manage the software and machines in an administration domain. This chapter explains how to use the different consoles in the module.

Topics

- [Resource Management Overview, page 107](#)
- [Working With Application Domains, page 108](#)
- [Customizing the Installed Software Display, page 111](#)
- [Customizing the Machines Display, page 112](#)
- [Disabling and Enabling Installed Software, page 113](#)
- [Adding Custom Software, page 114](#)
- [Configuring Monitoring for a Machine, page 116](#)
- [Removing a Machine from a Domain, page 118](#)
- [Turning Auto Refresh On or Off, page 119](#)
- [New Application Domain Panel, page 120](#)
- [Installed Software Dialog, page 122](#)
- [Machines Dialog, page 125](#)
- [View Machine Dialog, page 127](#)
- [Add Event Dialog, page 130](#)

Resource Management Overview

The Resource Management module allows you to manage application domains, view all installed software or product components registered in the administration domain, and manage machines in the domain. You can:

- Create application domains.
- Customize the machines display.
- Display information about the machine and processes running on the machine.
- Specify monitoring options for each machine in the domain. You can specify alert events or TIBCO Hawk rulebases. See [Configuring Monitoring for a Machine on page 116](#).



By default, if using Rendezvous for the domain transport, all machines within an administration domain are expected to be in the same network subnet. You can, however, set up your system to use TIBCO Rendezvous remote daemon (`rvrd`) to use TIBCO Administrator across subnets. See the *TIBCO Administrator Server Configuration Guide* for details. This does not apply for EMS domains.

Working With Application Domains

An application domain stores application data in a repository that is separate from the administration domain repository. This allows you to store configuration information required by the application in a repository that is independent of the administration domain repository.

An application domain is configured under an administration domain and uses the storage type defined for the administration domain. If an administration domain uses a file-based repository, each application domain also uses a file-based repository. If an administration domain uses a database repository, each application domain also uses a database repository. Multiple application domains can be configured under the same administration domain.

The domain administrator for the administration domain is also the domain administrator for each application domain created inside the administration domain.



This feature is available only if your TIBCO product supports it. If a TIBCO product does not support this feature, such as BusinessWorks or TIBCO Adapters, you can create application domains, but will be unable to assign an application to an application domain.

By defining an application domain, you can specify that one or more applications use a different repository to store application generated data.

For example:

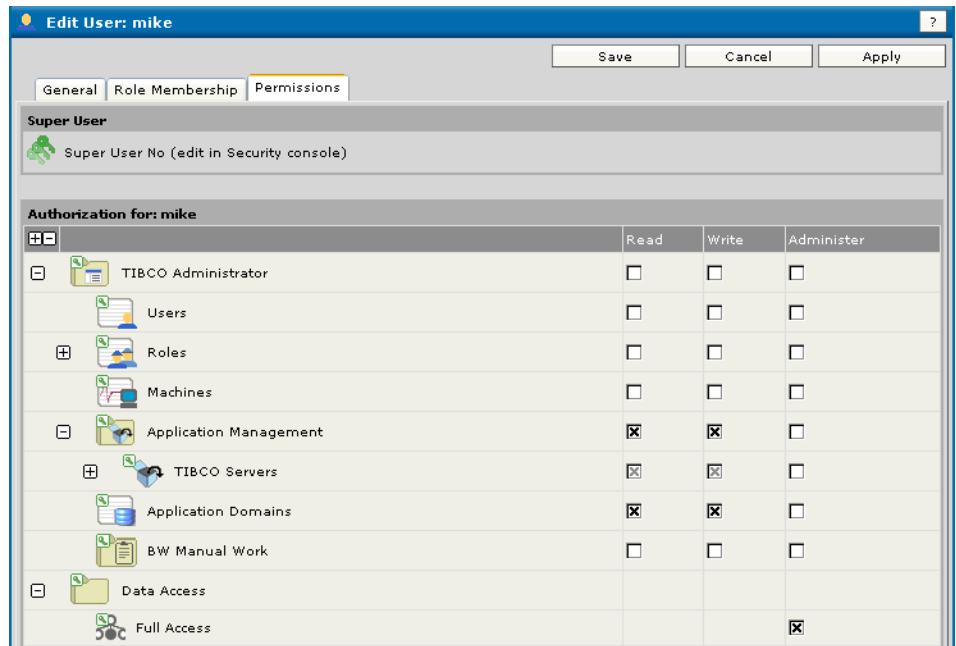
- Applications can use local databases for repositories instead of storing application data in the database used by the administration domain.
- If you are deploying multiple applications that are based on different versions and there is a model change between these versions, you can select a separate application domain for each version.
- If you are using a file-based repository for an administration domain and the repository is large in size, you can separate application data into smaller repositories.

You must first create an application domain and then assign your applications to the domain. See your TIBCO product's documentation set for more information.

Permissions

A user must have Write permission set for the Application Domains console. If the administration domain uses a file-based repository, the user must also have Data Access permission set as shown in the next diagram.

Figure 28 Setting DATA Access Permission



To Create an Application Domain

1. Click **Resource Management > Application Domains**.
2. Click **New**.
3. Provide values in each field. See [New Application Domain Panel on page 120](#) for field descriptions.
4. Click **Save**.



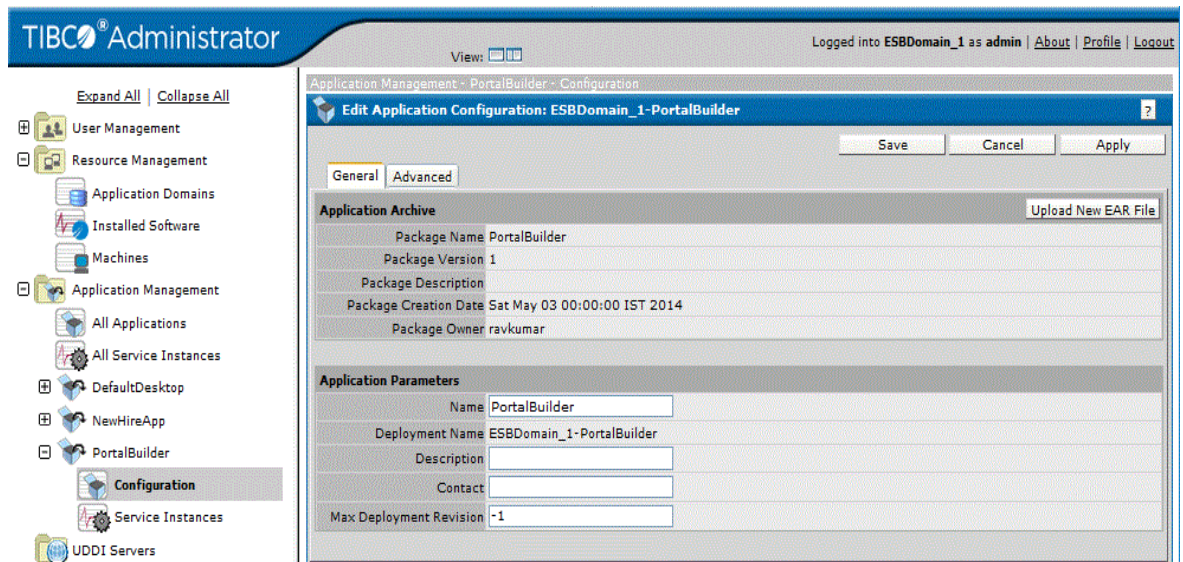
If the domain database password has been changed, confirm that the Application Domain is updated to reflect this change.

To Assign an Application to an Application Domain

This example uses an application that is part TIBCO PortalBuilder.

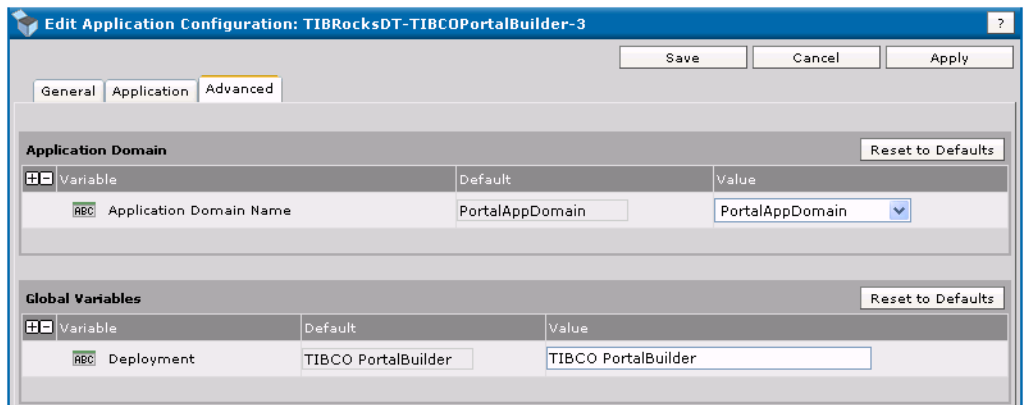
1. Expand **Application Management** and expand the application.
2. Click **Configuration**.
3. In the Configuration Builder panel, click the application name.

Figure 29 Assigning an Application to an Application Domain



4. In the Edit Application Configuration panel, click the **Advanced** tab.
5. Under Application Domain, click the Value field arrow and select the application domain to associate with the application.

Figure 30 Advanced Tab



6. Click **Save**.

Customizing the Installed Software Display

You can customize the Installed Software console by selecting a choice in the Group by combo box:

- Choose Machine to display installed software grouped by the machines in the domain.
- Choose Software to display grouped by installed software.
- Choose Enabled to view only software that is active in the administration domain.

By default, all installed software is displayed. To view only some of the components, type a search string into the search box, and then click **Search**. You can use the * character as a wildcard.

Customizing the Machines Display

You can customize the Machines console by selecting Group by: Hawk Cluster or Group by: OS/Version and:

- Click the minus(-) box to show only the machine group, but not the machines in it.
- Click the plus (+) box to show all machines and information about them for that machine group.

By default, all machines in all machine groups in the administration domain are displayed. To view only some of the machines, type a search string into the search box, and then click **Search**. Machine status is only available for machines on which the TIBCO Hawk agent service for the domain is currently running

If Disk Information Does Not Display

If the disk information for any machine running Microsoft Windows is not shown in the TIBCO Administrator GUI display, check the following:

1. Is the TIBCO Hawk agent service running on the remote machine?
2. If yes, run `diskperf` (with no option) to check if all the disk performance counters are enabled.
3. If the counters are not enabled, run `diskperf -y`, to enable all the counters. You need to execute this command only once; it will be executed automatically the next time you reboot your machine.



Disk performance counters are permanently enabled on Microsoft Windows 2000 and later Microsoft Windows versions.

Disabling and Enabling Installed Software

The Installed Software console allows you to disable and enable installed software. When a machine is part of an administration domain, all software from TIBCO becomes known to the corresponding administration server and visible in the TIBCO Administrator GUI. This will automatically update as long as the TIBCO Hawk is running on the machine that hosts the software. If the TIBCO Hawk agent is stopped, when restarted, it checks for TIBCO software on that machine and updates the information the administration server has about the software in the domain.

When you disable software on a machine, that software on that machine is no longer available for deploying applications. For example, if you installed TIBCO BusinessWorks software on two machines, you can disable the software on one machine to force any deployment to go to the other machine.

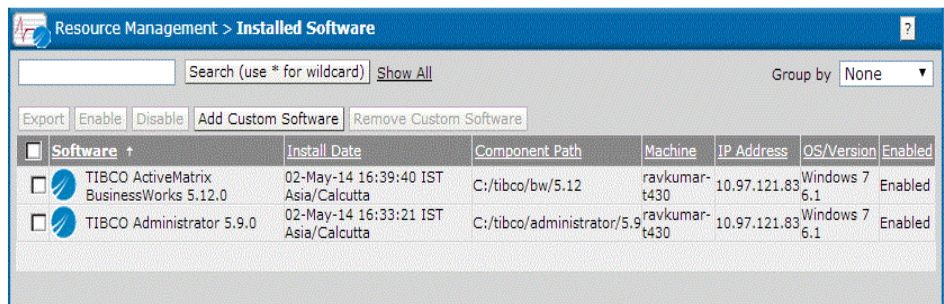
Enabling software is usually done after you have disabled it and is not necessary otherwise.

To Disable Installed Software

1. Click **Resource Management > Installed Software**.
2. In the Installed Software console, select the software to disable.
3. Click **Disable**.

The software is displayed with a red X over it.

Figure 31 Disable Installed Software



To enable disabled software, select it and click **Enable**.

Adding Custom Software

You can manage custom software from the Installed Software console. Once the software has been added to the domain, you can deploy applications that use that software just as you would deploy any other TIBCO application.

When a TIBCO product such as a TIBCO adapter or TIBCO BusinessWorks is installed on a machine that belongs to an administration domain, the software is automatically added to the domain. Similarly, if a machine is added to an administration domain, the TIBCO products on it are automatically added to the domain. However, custom software, such as adapters built using the TIBCO Adapter SDK must be added explicitly.

You can remove custom software from the display by selecting the software and clicking **Remove**.

To Add Custom Software

1. Click **Resource Management > Installed Software**.
2. Click **Add Custom Software**.
3. Select the machine to which you wish to add the custom software, and then click **OK**.
4. Provide information in the panel that is displayed, and then click **OK**.

See Also

For field descriptions, see [Installed Software Dialog on page 122](#).

If the machine on which the custom software is installed is not part of the administration domain, you must first add the machine to the domain. See the *TIBCO Runtime Agent Domain Utility User's Guide* for details.

Exporting Inventory Information to a File

You can export software inventory information to a .csv (comma-separated values) file, which can be opened with Microsoft Excel under Microsoft Windows and with other editors or spreadsheets on other platforms.



This functionality is not available in conjunction with Microsoft Internet Explorer 5.5. It is available with Internet Explorer 6.x or with Firefox 7.x.

To Export Inventory Information

1. Click **Resource Management > Installed Software**.
2. Select the items you want to include.
3. Click **Export**.
4. Supply the requested information.

Configuring Monitoring for a Machine

You can specify when and how you wish to be alerted by TIBCO Administrator if problems occur. You can either use a preconfigured TIBCO Hawk rulebase or configure alert events manually.

A Hawk rulebase is created using the TIBCO Hawk product. TIBCO Hawk allows you to specify a very large number of alert conditions and alert results. You must have purchased the full TIBCO Hawk product to configure and use custom TIBCO Hawk rulebases. A TIBCO Hawk agent service must run on the machine to which you wish to send the alert if you wish to create, modify, or delete alerts.

See [Adding a Custom Rulebase to a Process or Service on page 170](#) for information about creating a Hawk rulebase.



Removing a rulebase from the list and then sending the updated list of rulebases to the machine will not cause the removed rulebase to be cleared. You must explicitly clear the monitoring configuration first and then send the updated monitoring information using the **Send Configuration to Machine** button. However, for modifying existing rulebases, you can just send without clearing.

If you choose to specify an event, you can only select from the subset of the TIBCO Hawk conditions and results that are fully integrated with TIBCO Administrator.



After Domain is created, the hawkagent.cfg file of the created domain under TIBCO_Home/tra/domain/domain_name should enable SMPT settings, for sending an email alert or else an email action does not work. Then it shows an error in Hawk.log of the respective domain. Refer to, [Required Configuration for sending an email, on page 131](#) for details.

To Add a PreConfigured TIBCO Hawk Rulebase

1. Click **Resource Management > Machines**.
2. Click a machine name.
3. Click **Configure Monitoring**.
4. Click **Edit**.
5. Under Rulebases, click **Add**.
6. Click **Browse** and navigate to the directory where your TIBCO Hawk rulebases are defined. Select a rulebase and click **Open**.
7. Click **OK**.

8. Click **Done**.

To Add an Event

1. Click **Resource Management > Machines**.
2. Click a machine name.
3. Click **Configure Monitoring**.
4. Click **Edit**.
5. Under Events, click **Add**.
6. Define the event that should trigger the alert. Depending on the event type, the rest of the event panel will change.
7. Click **OK**, and then click **OK** in the next screen.
8. Click **Send Configuration to Machine**.

To Clear a Monitoring Configuration From a Machine

1. Click **Resource Management > Machines**.
2. Click a machine name.
3. Click **Configure Monitoring**.
4. Click **Clear Configuration from Machine**.
5. Click **OK** to confirm the message.
6. Click **Done**.

The alerts are now deactivated, but are still stored on your machine for later use if needed. Click **Send Configuration to Machine** if you wish to reactivate the alerts.

To Delete a Monitoring Configuration From a Machine

1. Click **Resource Management > Machines**.
2. Click a machine name.
3. Click **Configure Monitoring**.
4. Click **Edit**.
5. Under Rulebases or Events, click the configuration to remove.
6. Click **Delete**.
7. Click **OK**.
8. Click **Done**.

Removing a Machine from a Domain

You must stop and remove the process engines or service instances running on a machine before removing the machine from an administration domain. You are prompted to clean up dependencies if you do not perform the task.

To Stop and Remove Services

1. Click **Application Management**.
2. Choose **All Service Instances**.
3. Select the process engines or service instances running on that machine.
4. Click **Stop**.
5. In the Configuration Builder panel for the application, select the process engine (.par) or service instance (.aar).
6. With the General tab selected, click any service instances that were added to that machine in the Target Machines pane.
7. Click **Remove From Selected Machines**.

To Remove a Machine From a Domain

1. Click **Resource Management**.
2. Select **Machines**.
3. In the Machines console, click the box next to the machine name.
The machine is highlighted.
4. Click **Remove**.
5. Click **OK** in the confirmation dialog to remove the machine from the administration domain.

Turning Auto Refresh On or Off

TIBCO Administrator allows updates from events, rulebases or alerts to occur in real-time. Dialogs that use this feature include an auto refresh icon as shown in the next diagram. When auto refresh is enabled, the icon is active and displays movement across the horizontal bar.

Each time you click the auto refresh icon, auto refresh toggles between on and off.

Figure 32 Auto Refresh Icon



Auto refresh icon

You can set the amount of time that can expire before refreshing. See [Profile Dialog on page 93](#) for details.

New Application Domain Panel

See [Working With Application Domains on page 108](#) for more information.

Domain Name

The name for the application domain.

Domain Type

Lists the storage type to use for the application domain, either **Repo** or **Database**. The domain type is determined by the storage type used by the administration domain.

JDBC Driver

Java Database Connectivity (JDBC) drivers enable an application to connect with a database. Select from the following drivers:

- tibcosoftwareinc.jdbc.oracle.OracleDriver
- tibcosoftwareinc.jdbc.sqlserver.SQLServerDriver
- weblogic.jdbc.mssqlserver4.Driver

Database DSN

The database DSN is required by the driver. It identifies which database is used for the session.

Database Username and Database Password

Enter the name and password required to connect to the database. The database user must have permission to:

- Create and drop tables, indexes and sequences.
- Insert, delete, and update tables.

Min Connections

When your application starts up, it initializes the number of connections to the database specified in this field. For optimal performance, this number should equal the number of subscribers you expect to connect to the database at any one time.

As you specify the minimum connections, keep in mind that these are per instance, and at startup, each instance will open its specified number of connections. Higher settings are better for application instances, but may have adverse results in the rest of the system. The correct setting is whatever the complete system can handle on a sustained basis without taxing other system resources.

Max Connections

The number of simultaneous connections cannot exceed the number set in this field. The database must be able to simultaneously handle the total maximum number of connections for all instances.

Be sure that you set Max Connections to a number greater than the number for Min Connections. Otherwise you will get an error message.

TIBCO recommends that you set Max Connection to a number equal or greater than the default value. TIBCO Administrator server needs a suitable amount of connections to initialize and to deal with domain data stored cross multiple tables in the domain. If the value of Max Connection is less than the default value, the operations may start up slowly and a timeout error may occur.

Table Prefix

The table prefix that is appended to each table used by the application. Can be no more than two characters.

Test Connection

Click this button to verify that the values you have provided result in a valid database connection.

Installed Software Dialog

Search

Allows you to display only the items that match a search criteria. You can use the * character as a wildcard.

Group By

Determines how the items in the display are grouped. Select from the following choices:

- None — Do not use grouping.
- Machine — List first all components on one machine, and then all components on the next machine in the administration domain.
- Software — Organize the display by first showing all components (service instances and process engines) then showing the machine.
- Enabled — Group by enabled and disabled components.

Export

Allows you to export inventory information for the selected item(s) to a .csv (comma-separated values) file, which can be opened with Microsoft Excel under Microsoft Windows and with other editors or spreadsheets on other platforms.

Enable, Disable

When a machine is part of an administration domain, all software from TIBCO becomes known to the corresponding administration server and visible in TIBCO Administrator updates occur automatically while TIBCO Hawk agent service is running on that machine. If the TIBCO Hawk agent is stopped, it checks for TIBCO software on that machine when it is restarted, and updates the information the administration server has about the software in the domain.

When you disable software on a machine, that software on that machine is no longer available to applications. As a result, the software is not available in the deployment configuration dialog. For example, if you installed TIBCO BusinessWorks on two machines, you can disable the software on one machine to force any deployment to go to the other machine.

Enabling software is usually done after you have disabled it and is not necessary otherwise.

Add Custom Software

Allows you to add custom software to this administration domain. After selecting the machine on which the software is to be added and clicking OK, the following fields become available:

- **Machine** — Name of the machine to which the software is to be added. Click **Change** to add to a different machine.
- **Software Type** — Must match the software type used to build an enterprise archive file.
- **Software Display Name** — Named displayed in when monitoring the software.
- **Version** — A four-component number indicating the software version, for example 5.1.2.9. This number is then matched against enterprise archive files loaded for this software. If the enterprise archive file specifies a later version, it cannot be deployed.
- **Executable (Full Path)** — The full path to the executable for this custom software.
- **Software is an adapter** — Select this box if your custom software is an adapter, clear the box otherwise.
- **Java Software** — Select this box if your custom software is written in Java, clear the box otherwise.
- **Java Start Class** — Provide the Java start class.
- **Java Start Method** — Provide the Java start method.
- **Java Stop Method** — Provide the Java stop method.
- **Java Classpath** — Provide the Java classpath.

Remove Custom Software

Select the custom software package to remove from this list and click this button.

Software List

- **Software** — Name of the installed software.
- **Install Date** — Time and date at which this component was installed.
- **Component Path** — Location of the component on the machine displayed in the Machine column.
- **Machine** — Name of the machine on which the component has been installed.
- **IP Address** — IP address of the machine.

- OS/Version — Operating system information.
- Enabled — Displays enabled if this software has been enabled (the default), Disabled otherwise.

Machines Dialog

Search

Allows you to display only the items that match a search criteria. You can use the * character as a wildcard.

Remove

Removes the selected machine from the domain.

Group By

Determines how the display is grouped.

- None — Do not use grouping.
- Hawk Cluster — List first all machines in one group, and then all machines in another.
- OS/Version — List first all machines that belong to one operating system, and then all machines that belong to another.

Machine List

- Hawk Cluster — Name of the group in which this machine is running. Default is the network IP address of the group.
- Machine — Click any machine’s name for additional information about that machine. See [View Machine Dialog on page 127](#) for details.
- Status — Current status for this machine. One of following is displayed:

Table 4 Status Icons






Icon	Description
	Status is normal.
	Indicates alert severity is set to low.
	Indicates alert severity is set to medium.

Table 4 Status Icons

Icon	Description
	Indicates alert severity is set to high.
	Indicates that contact to the endpoint Hawk Agent is lost. The network may be down, or the machine on which the Hawk Agent is running may be off line.

- Uptime — Displays the number of days and hours this machine has been running since it was added to TIBCO Administrator, or since the machine has been rebooted. It requires the TIBCO Hawk agent to be running on that machine.
- OS/Version — Operating system and version running on this machine.

View Machine Dialog

The following tabs are available:

- [General Tab](#)
- [Processes Tab](#)
- [Configure Monitoring Tab](#)

General Tab

General Information

Displays general information about the machine including:

- Status — State of this machine
- Uptime — Displays the number of days and hours this machine has been running since the machine has been rebooted.
- IP Address — IP address of this machine
- OS/Version — Operating system and version.
- Virtual Memory Usage — The virtual memory usage for this machine. If any bars are yellow or red, there might be a problem.



"Memory usage (in KB) displayed on Tibco Admin GUI represents the virtual memory consumption by a running process on UNIX/Linux platform which indicates how much memory the program is able to access at present. This is not the actual physical memory a process consumes.

While on the windows platform, memory usage (in KB) displayed represents the working set memory. The working set of a process is a collection of those pages in its virtual address space that have been recently referenced. It includes both shared and private data."

- CPU Usage — CPU usage for that machine. There is one bar for each CPU. Bars change to yellow or red for high usage.
- Disk Usage — Disk usage on that machine. There is one bar for each disk on the machine. Bars change to yellow or red for high usage.



It requires TIBCO Hawk agent to be running on the machine.

Active Alerts

Shows all currently active alerts on this machine. Alerts can be configured using the Configure Monitoring tab. See [Configuring Monitoring for a Machine on page 116](#) for information on configuring alerts.

- Date/Time — Date and time of this alert.
- Alert Level — Depending on whether you chose High, Medium, or Low when you configured this alert, the display will change.
- Source — Alert source, such as Machine, Software Container, or Process.
- Text — Alert text.

Processes Tab

When you click the Processes tab, the display changes to show processes running on that machine. By default, only TIBCO processes are displayed. You can click **Show All Processes** for a display of all processes running on that machine; then click **Hide Non-TIBCO Processes** to show only TIBCO processes (the default).

Show All Processes

The following information is available in the Processes display:

- PID — Process identifier.
- Process name — Name of the process.
- CPU (%) — Percent of CPU this process is currently consuming.
- Memory (KB) — Number of kilobytes used by this process.
- TIBCO Components — Product name of the deployed application or TIBCO service.



It requires TIBCO Hawk agent to be running on the machine.

Configure Monitoring Tab

Clear Configuration from Machine

Clears a monitoring configuration from a machine. The alerts are now deactivated, but are still stored on your machine for later use if needed. Click **Send Configuration to Machine** if you wish to reactivate the alerts

Send Configuration to Machine

Sends a monitoring configuration you prepared to the local machine. A TIBCO Hawk agent must be running on the machine to which you wish to send the alert if you wish to create, modify, or delete alerts.

Add Event Dialog

The following sections are available:

- [General Pane](#)
- [Alert Pane](#)
- [To send an email, we have to configure few settings in TIBCO_Home/tra/domain/domain_name/hawkagent.cfg specified below.](#)
- [Command Pane](#)

General Pane

Depending on the event type selected and whether you are setting an event for a machine or a process, the fields on the event panel change.



If you are setting an event for a machine, the following fields display:

Event type CPU

- CPU — For multi-CPU machines, select the CPU you wish to monitor.
- CPU Usage — Choose first the operator, and then the number. For example, you may want an alert if CPU Usage is greater than or equal to (\geq) 80%.
- Description — Optionally, provide additional information, which will be picked up by the alert itself.

Event type disk

- Disk — The disk on this machine that you wish to monitor.
- Disk Free Space — Choose first the sign, and then the number. For example, you may want an alert if disk free space is less than or equal to (\leq) 90%.
- Description — Optionally, provide additional information that will be picked up by the alert itself.

Event type process

- Process Count — Choose first the operator, and then the number. For example, you may want an alert if there are more than fifty (>50) processes running.

Description

Optionally provide additional information that is picked up by the alert itself.

Alert Pane**Generate Alert**

Select this box if you wish to generate an alert to be sent to TIBCO Administrator and the log.

Generate Alerts for

Specify whether you wish to generate an alert for the first occurrence only or for all occurrences.

Level

Choose one of High, Medium, or Low. This will affect the appearance of the alert in TIBCO Administrator.

Message

Provide the message to display when the alert is triggered.

Required Configuration for sending an email

To send an email, we have to configure few settings in TIBCO_Home/tra/domain/domain_name/hawkagent.cfg specified below.

Here, SMTP HostName is the mandatory field.

```
# Email Action configuration
# SMTP server host name
#-email_smtp_server <SMTP HostName>
```

You can change default values.

```
# SMTP server port, default value is 25.
#-email_smtp_port 25
```

```
# Whether SMTP server authentication is required or not,
default value is false.
```

```

#-email_smtp_auth_required false

# SMTP server user name
# Required only if SMTP server authentication is configured
true.
#-email_smtp_user <UserName>

# SMTP server user password.
# Required only if SMTP server authentication is configured
true.
#-email_smtp_password <password>

# Email From (Email Sender)
# Optional, default is current system user
#-email_from "HawkAdministrator"<admin@abc.com>
#-email_from <Sender email address>

```

It is mandatory to restart Hawkagent after the changes.

Email Pane

Send Email

Select this box if you wish for an email to be sent to a specified user.

Send Email For

Specify whether you wish to send an email for the first occurrence only or for all occurrences.

Mail Server

Provide the mail server (SMTP server) to use to send the message. You can specify the host name or the host IP address.

To

Provide a comma-separated list of email addresses to which you would like to send the message.

CC

Provide a comma-separated list of email addresses to which you would like to send copies of the message.

Subject

Provide the subject of the email message.

Message

Provide the text of the email message.

Command Pane**Execute Command**

Select this box to execute a command when the specified condition is met.

Execute Command For

Specify whether you wish to execute a command for the first occurrence only or for all occurrences.

Command

Specify the script to execute. Script files are highly recommended.

Commands are possible but are limited because the command line arguments cannot accept redirection (`|`), multiple command (`;`) or append characters (`>` and `>>`). Redirection is allowed in a script.

On Windows:

- Use a `.bat` file that begins with the line `@echo off` to prevent the shell from exiting prematurely.
- Always give the full path with `"\"` as the path separator.
- If you use a command instead of a script, you must prefix it with `cmd \c.`

On UNIX:

- Make sure the script is executable (`chmod +x`).
- Always give the full path with `/` as the path separator.

If you purchased the full TIBCO Hawk product, see the *TIBCO Hawk Administrator's Guide* for more information.

Arguments

To pass to the script, specify the list of arguments.

Rulebases

Lists the rulebases defined on this machine.

Events

Lists the event types defined on this machine.

Edit

Edits a monitoring configuration.

Click Add to display the Upload Rulebase File dialog where you can browse the file system for an existing TIBCO Hawk rulebase file. See To Add a Custom Rulebase to an Application on page 153 for more information.

Click Add to display the Add Event dialog. See Add Event Dialog on page 114 for details.

Chapter 7 **Creating and Deploying Applications**

The TIBCO Administrator Application Management module allows you to create and deploy applications using the corresponding archive files.

Topics

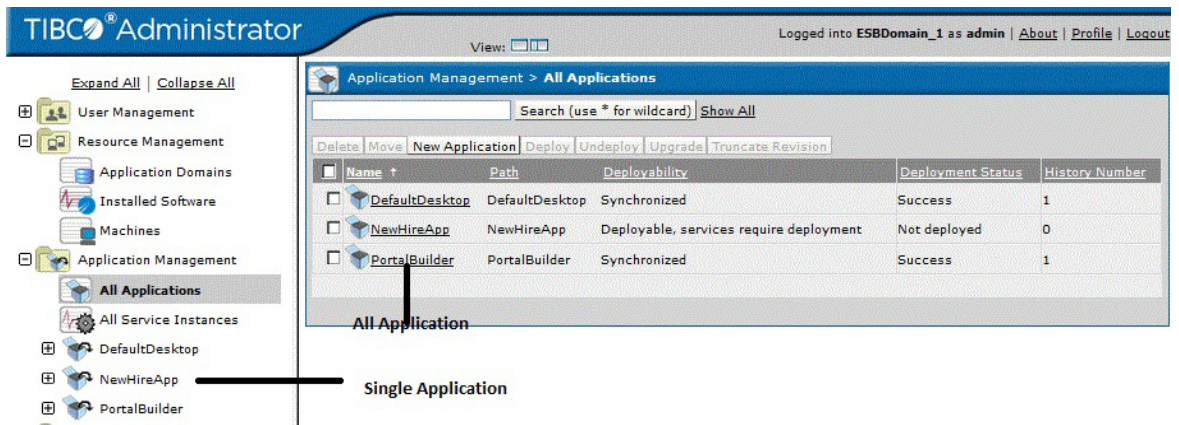
- [Application Management Overview, page 136](#)
- [Creating an Application, page 137](#)
- [Deleting an Application, page 140](#)
- [Deploying an Application, page 141](#)
- [Reverting to a Previously Deployed Application, page 144](#)
- [Undeploying an Application, page 145](#)
- [Deploying an Application Using Dynamic Encryption Key, page 146](#)
- [Managing Folders, page 147](#)
- [Viewing Application Deployment History, page 149](#)
- [Purging Application Revisions, page 150](#)
- [Upgrading an Application, page 151](#)
- [Application Management Dialog, page 154](#)
- [All Applications Dialog, page 155](#)
- [New Application Configuration Dialog, page 156](#)
- [Deploy Configuration Dialog, page 159](#)
- [View History Dialog, page 161](#)

Application Management Overview

The Application Management module allows you to create and deploy applications, and then start, stop, and monitor them.

The module contains the application you have loaded into TIBCO Administrator. You can view all applications in the All Applications dialog or all service instances and process engines in the All Services dialog. Alternatively, you can configure and manage the service instances and process engines for an application under the application's dialog.

Figure 33 Application Management



Creating an Application

To create an application in TIBCO Administrator, you must import an enterprise archive file, which is created in TIBCO Designer. An archive file typically contains a project that includes one or more configured adapter services, TIBCO BusinessWorks processes, or both.

After creating the application, you specify deployment configuration information, such as which machines should run which services and processes in the application. You then deploy and start the services and processes from TIBCO Administrator.

You can use the same enterprise archive file to create multiple applications and configure and deploy each application separately with different deployment options. You can also modify a deployed application and redeploy, and then revert to an earlier deployment if the changes you made do not have the desired result.

To Create an Application

1. Select either **Application Management**, a previously created folder or **All Applications**.
2. Click **New Application**.
3. Click **Browse** and select an enterprise archive file, and then click **OK**. For example, the next diagram shows a new application, ready for deployment.



It is recommended to use alphanumeric characters, dash (-), and underscore(_) for the application names and deployment names.

Figure 34 Creating an Application

Application Management - All Applications

New Application Configuration: FileMonitor

Save Cancel

General

Application Archive [Change EAR File](#)

Package Name: FileMonitor

Package Version: 20

Package Description:

Package Creation Date: Mon Jan 21 14:24:00 PST 2008

Package Owner: huchangcheng

Application Parameters

Name: FileMonitor

Deployment Name: sample-FileMonitor

Description:

Contact:

Max Deployment Revision: -1

Services

Quick Configure: ☒

Deploy on Save: ☐

Service	Description	Target
Process1.par		clin - new bwengine 5.6.0.7

4. Click **Save**.

Application Creation Choices

You have the following choices when creating an application:

- You can verify application information and make choices in the fields that allow input. If you wish, you can select a different archive file by clicking **Change EAR File**.
- If **Quick Configure** is selected, the services are bound to the targets selected in the target field.
- If **Quick Configure** is selected, **Deploy on Save** can be selected. When the **Save** button is clicked, the application is created and immediately deployed to the target machines specified in the Services pane, Target column. All variables, logging and other configuration values will use defaults defined in the archive file. The next screen will display the deployment status.


If **Deploy on Save** is not selected, the application must be explicitly deployed using the application's Configuration Console that displays upon save. This

allows you to change settings, such as global variable settings before deploying.

Application Types

Some TIBCO applications provide EAR files that are not created using TIBCO Designer. These files are known as application types and are installed when the application is installed. If your application provides these EAR files, a dialog similar to the following appears when creating an application. Pick the application type to load and click **OK**.

Figure 35 Uploading EAR File



After you load an application type, it displays as an application under the Application Management module.²

Deleting an Application

When you delete an application, all files associated with that application are removed and it becomes unavailable from TIBCO Administrator. You must upload the related enterprise archive file again to recreate the application if you wish to use it.

A deployed application must be undeployed before it can be deleted.

To Delete an Application

1. Click **Application Management** or **All Applications**.
2. Select the application to delete.
3. Click **Delete**.
4. Click **OK** in the confirmation dialog.

Deploying an Application

When you create an application, you can use the Deploy on Save option to deploy the application when saving. The application then uses the global variables and other options set in the enterprise archive file. If you do not use the Deploy on Save option, you must explicitly deploy the application after changing deployment configuration options.

To Deploy an Application

1. Click **Application Management**.
2. Click *Application* > **Configuration**, where *Application* is the application created when you loaded the enterprise archive file. See [Creating an Application on page 137](#).
3. Before deploying, you can change the deployment options. See [Deployment Choices on page 142](#) for details.
4. Click **Deploy**.

The dialog similar to the following displays and informs you that all running processes with configuration changes in this application will be stopped when you click **OK** to deploy. If the processes deploy successfully, they are restarted automatically if the corresponding check box was selected.

Figure 36 Deploying an Application

Application Management > FileMonitor > Configuration

Deploy Configuration

OK Cancel

Stop running services before deployment	<input checked="" type="checkbox"/>
Kill services that haven't stopped after (seconds)	5
Start successfully deployed services	<input checked="" type="checkbox"/>
Force redeployment of all services	<input type="checkbox"/>
Dynamic Symmetric Key	<input type="checkbox"/>
Description	

☒ **FileMonitor**

Administrator Tasks To Perform

Create/Update the TIBCO Repository instance that the BwEngine or Adapters will use.

Remote Tasks To Perform	Service Instance	Service Configuration	Deployability
Deploy Engine on Client	clin - Process1	Process1.par	Deployable, (New)

The dialog allows you to add a description and displays information about the application and each service.

5. Click **OK** if to deploy the application, or click **Cancel** to choose an advanced configuration, different archive file, or make other changes.

Deployment Choices

You can make changes to a deployed application, and then deploy the changed application. The currently deployed application can continue to run while you make changes. When you deploy the updated application, the current application is automatically undeployed. You can revert to a previously deployed application, if the changes you made need be rolled back.

You can:

- Select **Stop running services before deployment** to stop all running services before deploying the service. All services that should be redeployed are stopped. If, however, a service instance is in sync and **Force redeployment of all services** is not enabled, the running instance is not stopped. This option is disabled if any target machine has a TIBCO Hawk agent version lower than 5.3.
- Indicate how many seconds can elapse after a stop request before a service is killed, using the **Stop** command, in **Kill services that haven't stopped after (seconds)**.
- Select **Start successfully deployed services** to stop and restart the services in the application after they have been successfully deployed. If you do not select this option, you can explicitly start the services later. See [Starting or Stopping a Service Instance or Process Engine on page 215](#) for details.
- Select **Force redeployment of all services** to redeploy all services even if a service is in a synchronized state. This is useful if you have manually changed deployment files, or if you need to define NT Services to multiple hosts in a Microsoft Cluster.
- Select **Dynamic Symmetric Key** to encrypt sensitive application data using a dynamically generated encryption key. See [Deploying an Application Using Dynamic Encryption Key on page 146](#) for more information.

If the **Stop running services before deployment** option is cleared, service instances in deployable state are not stopped while they are redeployed. However, they will be stopped after a successful deployment. This option has no impact on the services instances that are in a synchronized state unless the **Force redeployment of all services** option is selected.

If the **Start successfully deployed services** option is enabled and the **Stop running services before deployment** option is cleared:

- Existing service instances in running state are first stopped then restarted after .tra and .cmd files on target machines are updated
- Existing service instances not in running state are restarted when they are deployed.
- Newly added service instances are restarted.
- Removed service instances are stopped before they are removed.

Reverting to a Previously Deployed Application

When you revert an application, you select a different version of the currently deployed application to deploy. When you deploy, service instances and process engines are stopped, updated, and restarted. Any component that is removed from a machine as a result of the revert operation is undeployed from that machine.




Reverting an application is only possible if you have deployed an application more than once.

To Revert to a Previously Deployed Application

1. Click **Application Management**.
2. Click *Application* > **Configuration**, where *Application* is the application created when you loaded the archive. See [Creating an Application on page 137](#).
3. In the Configuration Builder panel, click **Revert**.

The deployed revisions and the time at which each was deployed display.

Figure 37 Reverting to a Previously Deployed Application



4. Click the button next to the revision you wish to use.
5. Click **OK**.

The application is now shown as ready to deploy in the Configuration Builder.

6. Click **Deploy**.

Undeploying an Application

When you undeploy an application, TIBCO Administrator stops all running services and removes them from the list of services that can be started. In effect, it completely removes all traces of the deployment (with the exception of the logs).



To use an earlier version of the deployment configuration, select **Revert**, not **Undeploy**. See [Reverting to a Previously Deployed Application on page 144](#).

To Undeploy an Application

1. Click **Application Management**.
2. Click *Application* > **Configuration**, where *Application* is the application created when you loaded the archive.
3. In the Configuration panel, click **Undeploy**.
4. Click **OK** to undeploy, or **Cancel** to stop the operation.

Undeploy Dialog

Kill services that haven't stopped after (seconds)

Specify the amount of time to wait before killing service instances or process engines that have not stopped. The default is zero, meaning no time is allowed for a graceful shutdown, even if graceful shutdown had been previously.

Administrator Tasks To Perform

Lists the tasks that TIBCO Administrator will perform for this server if you choose to undeploy by selecting the OK button.

Remote Tasks To Perform

Lists the tasks to perform on the selected machine (which could actually be the local machine) in the following fields:

- Service Instance.
- Service Configuration.
- Deployability — Shows whether the application has been deployed before.
- Task — Actions the deployment process performs on the target machine(s).

Deploying an Application Using Dynamic Encryption Key

When deploying applications on client machines, TIBCO Administrator encrypts all the sensitive data in the applications, such as password-type variables. While TIBCO Administrator uses a static encryption key for such encryptions, you can also use a dynamically generated encryption key to provide stronger security for your sensitive data on the client machines.

Note that the dynamic encryption key option is only available for domains stored in databases and that store application data on client machines, including domains that use TIBCO Enterprise Message Service. In other words, you must have selected both of the following checkboxes in TIBCO Domain Utility when creating your domain:

- Domain information stored in a Database
- Local Application Data



If the domain is not set to Local Application Data by default, you can still override it on a per application basis to use this feature.

See *TIBCO Runtime Agent Domain Utility User's Guide* for more information.

You can also encrypt application configuration files when exporting them using the AppManage utility. See *TIBCO Runtime Agent Scripting Deployment Guide* for more information.

To Use Dynamic Encryption Key

When deploying or redeploying an application, select the **Dynamic Symmetric Key** checkbox. When using this option, the communication between the TIBCO Hawk agent and the database will need to use JDBC over SSL to be secure.



If using this option, the .tra file and application data (.dat) file will only work on the target machine as the encryption key is tied to the machine.

Managing Folders

If the structure of the applications you expect to manage using TIBCO Administrator is complex, you can organize the applications into folders. After creating a folder, you can add applications to the folder and can also create sub-folders to form a tree.



When you delete a folder, the folder contents are also deleted.

To move a folder, you must have Administer permissions on the source folder (including its contents) and the destination folder.

To Create a Folder

1. Select either **Application Management**, or a previously created folder.
2. Click **New Folder**.
3. Provide a folder name and, optionally, a description and contact.
4. Click **Save**.

To Delete a Folder

1. Select the folder's parent, either **Application Management**, or a previously created folder.
2. Select the folder to delete.
3. Click **Delete**.
4. Click OK in the confirmation dialog.

To create or delete a folder, add applications to a folder, or create sub-folders, you must have write permission to the folder.

To Move a Folder

1. Select the folder's parent, either **Application Management**, or a previously created folder.
2. Select the folder to move.
3. Click **Move**.

Moving an Application to a Folder

If you wish to organize your applications into folders, or need to move your application for other reasons, you can do so from the Application Management console.



To move an application, you must have Administer permissions on the application and the destination folder.

To Move an Application to a Folder

1. Click **Application Management** or **All Applications**.
2. Create folders if desired.
3. Select the application you wish to move.
4. Click **Move**.
5. You are prompted for the desired location of the application.
6. Click **Save** to make the change.

Viewing Application Deployment History

You can view a history of each time an application has been deployed. History includes the date when deployed, the user who deployed the application and a description, if given when the application was deployed. The Status Details panel provides information about the success or failure of each deployed process in the application.

To View Deployment History

7. Click **Application Management**.
8. Select an application
9. Click **Configuration**.
10. Click **History**.
11. Click **Details** for more information.

Purging Application Revisions

You can specify how many recent revision entries to keep in the revision history for an application. By default, TIBCO Administrator maintains the history of all application revisions.

To Specify Maximum Revisions for an Application

1. Click **Application Management**.
2. Click *Application* > **Configuration**, where *Application* is the application created when you loaded the archive. See [Creating an Application on page 137](#).
3. Click the application name in the Configuration Builder area.
4. Change the **Max Deployment Revision** field to the number of revisions you wish to keep. This is -1 by default, meaning all revisions are kept in the history.
5. Click **Save**, then click **OK**.

The next time you deploy a new revision of the application, TIBCO Administrator will purge the extra old revisions from the revision history. For example, if there were ten revisions in the revision history and you specified 5 in the Max Deployment Revision field and deploy a new revision of your application, you can see that only the five most recent revisions are kept in the revision history.

To Purge Revisions Immediately

Click the **Truncate Revision** button in the application's Configuration panel. TIBCO Administrator will purge old revisions immediately, according to the value you specify in the Max Deployment Revision field for the application.

Upgrading an Application

If you have installed a new version of TIBCO software on a machine that is part of your administration domain, and the software is used in one or multiple applications, you can use the Upgrade feature to enable the applications to use the upgraded software.



The All Applications console can be used to easily upgrade multiple applications.

The Upgrade feature remaps properties in the service instances and process engines properties files to use the new software targets. After you have upgraded, you must continue to use the upgraded software. That is, you cannot revert to using the previous software version.

You must redeploy your applications after upgrading. The **Upgrade** panel gives you the option to deploy after upgrading. However, if you wish to redeploy using the dynamic encryption key, you must manually redeploy using the Deploy Configuration dialog. See [Deploying an Application on page 141](#)

The next diagram shows the dialog that is displayed when upgrading software.

Figure 38 Upgrading a Software

Application Management > All Applications

Select Software to Upgrade

OK Cancel

Deployment Options

Deploy after upgrade ☒

Stop running services before deployment ☒

Kill services that haven't stopped after (seconds)

Start successfully deployed services ☒

Description

Select Software to Upgrade

<input checked="" type="checkbox"/>	Upgrade From	Upgrade To
<input checked="" type="checkbox"/>	TIBCO BusinessWorks 5.4.0	TIBCO ActiveMatrix BusinessWorks 5.6.0

Upgrade Summary

Name	Upgrade From	Upgrade To
sample-ActiveDatabase Adapter		
clin - Process1	TIBCO BusinessWorks 5.4.0	TIBCO ActiveMatrix BusinessWorks 5.6.0
sample-FileMonitor		
clin - Process1	TIBCO BusinessWorks 5.4.0	TIBCO ActiveMatrix BusinessWorks 5.6.0

To Upgrade an Application

1. Click **Application Management**.
2. Click *Application* > **Configuration**, where *Application* is the application created when you loaded the archive.
3. In the Configuration panel, click **Upgrade**.
4. Select **Deploy after upgrade** to redeploy your application as part of the upgrade. You can redeploy later.
5. Select **Start successfully deployed services** to deploy, stop and restart your service instances and process engines. If Deploy after upgrade is not selected, this option is not available.
6. Select the software to upgrade.
7. Review the upgrade summary.
8. Click **OK**.

See Also

See [Deploying an Application on page 141](#).

Application Management Dialog

Delete

Removes the selected applications from the domain. All deployment-related information is deleted. You must upload the related enterprise archive file again and create the application again if you wish to use it.

Move

Moves the selected application into another folder. You must have Administer permission for the application and folder.

New Folder

Click to create a new folder. After creating the folder, you can select it and add another folder, or an application.

New Application

Adds a new application to the domain. See [Creating an Application on page 137](#) for details.

All Applications Dialog

Use the All Application dialog to display all applications defined in the domain, provide a search criteria, and allow for actions against a number of applications at a time.

Delete

Removes the selected applications from the domain. All deployment-related information is deleted. You must upload the related enterprise archive file again and create the application again if you wish to use it. Applications must be deployed.

Move

Moves the selected application into another folder. You must have Administer permission for the application and folder.

New Application

Adds a new application to the domain. See [Creating an Application on page 137](#) for details.

Deploy

Deploy the selected applications. See [Deploying an Application on page 141](#) for details.

Undeploy

Undeploy the selected applications. See [Undeploying an Application on page 145](#) for details.

Upgrade

Upgrades the service instances and process engines in the selected applications to use a new software version. See [Upgrading an Application on page 151](#) for details.

Truncate Revision

Purge old revisions from the revision histories of the selected applications. See [Purging Application Revisions on page 150](#) for details.

New Application Configuration Dialog

The following sections are available:

- [Application Archive Pane](#)
- [Application Parameters Pane](#)
- [Services Pane](#)

Application Archive Pane

Displays information about the Enterprise Archive that was uploaded for this application.

Package Name

Name of the enterprise archive file, as specified in TIBCO Designer.

Package Version

Version of the enterprise archive file, as specified in TIBCO Designer.

Package Description

Description of the enterprise archive file, as specified in TIBCO Designer.

Package Creation Date

Date on which the enterprise archive file was created.

Package Owner

Owner of the enterprise archive file, as specified in TIBCO Designer.

Application Parameters Pane

Name

Name that this application will have in TIBCO Administrator. Defaults to the name of the enterprise archive file, but can be changed.

Deployment Name

Name of the application repository that is generated by TIBCO Administrator when you deploy the application. This file contains information about the application configuration. The file name defaults to the package name, with a number appended if more than one instance of the package has been deployed.



Do not manually edit this field.

Description

Description of this application.

Contact

Contact person for this application.

Max. Deployment Revision

The maximum number of revisions to keep in the revision history for this application. See [Purging Application Revisions on page 150](#) for more information.

Services Pane**Quick Configure**

This option allows you to quickly assign target machines to each service. If selected, the components will accept the defaults for global variables, options, and so on as set in the enterprise archive file. If cleared, components must be explicitly assigned to targets using the Configuration Builder for the application.

Deploy on Save

Available only if Quick Configure is selected.

If selected and Quick Configure is selected, and then, when the **Save** button is clicked, the application is created and immediately deployed to the target machines. All variables, logging and other configuration values will use defaults.

If this box is cleared, the application must be explicitly deployed using the application's Configuration Console.

Service, Description, Target

The Service column lists the services made available for this application. For example, if you included two TIBCO BusinessWorks process archives in your Enterprise Archive when you configured it in TIBCO Designer, each will be displayed as a service in this column.

The Description column lists the description specified for the service when defining the enterprise archive file.

The Target column offers a pop-up of potential service instances. For example, if you have TIBCO BusinessWorks on machine A and machine B, the pop-up for a TIBCO BusinessWorks service would show *A-process_name* and *B-process_name*. You can also choose to disable the service initially, and then enable it later.



If this column is empty for the service, TIBCO Administrator cannot find the associated software or it cannot find the correct version of the associated software. Check that the correct version of the associated software is on a machine in the TIBCO Administration domain. For custom adapters, you must also make sure the software is has been added to the administration domain. See [Adding Custom Software on page 114](#).

Deploy Configuration Dialog

See [Deployment Choices on page 142](#) for more information about these options.

Stop running services before deployment

Select **Stop running services before deployment** to stop all running services before deploying the service. All services that should be redeployed are stopped. If, however, a service instance is in sync and **Force redeployment of all services** is not enabled, the running instance is not stopped. This option is disabled if any target machine has a TIBCO Hawk agent version lower than 5.3.

Kill services that haven't stopped after (seconds)

Specify the amount of time to wait before killing services that have not shutdown.

Start successfully deployed services

If selected, TIBCO Administrator starts successfully deployed services. If the service is already running, it is stopped and then restarted. If cleared, services must be started explicitly.

Force redeployment of all services

If selected, all services are redeployed even if the service is in a synchronized state.

Dynamic Symmetric Key

If selected, TIBCO Administrator uses a dynamically generated encryption key to protect sensitive application data when deploying the application. See [Deploying an Application Using Dynamic Encryption Key on page 146](#).

Description

Provide a description for this deployment. This is useful for reviewing details and essential if you expect to revert later.

Administrator Tasks To Perform

Lists the tasks that TIBCO Administrator will perform for this server if you choose to deploy by selecting the OK button.

Remote Tasks To Perform

Lists the tasks to perform on the selected machine(s) (which could actually be the local machine) in the following fields:

- **Software** — Required software for this application (for example, an adapter or TIBCO BusinessWorks).
- **Deployability** — Shows whether the application is deployable and whether it's been deployed before.
- **Machine** — Computer on which the application is scheduled to be deployed.
- **Machine Status** — Machine state.
- **Tasks**— The actions that the deployment process will perform on the target machine(s).

Synchronized Tasks - process will continue running

Lists the processes that are currently part of the application and did not change in the new enterprise archive file that is being loaded and deployed. For example, if a TIBCO BusinessWorks project archive contains one process archive and you add a second project archive, the first project archive appears as a synchronized task unless you make changes to it.

View History Dialog

Allows you to view the deployment history.

Date

Displays the time when deployment occurred.

User

Displays the user who initiated deployment.

Action

Displays the action taken.

Revision Number

Indicates the number of times the deployment has been changed.

Description

Displays information about the deployment.

Details

When you click details, TIBCO Administrator displays a dialog with the following information:

- **Name** — Displays the components that must perform updates before this application can be started. This could be, for example TIBCO Administrator, the machine on which administrator is running, or the software.
For each component, you can view the change that was attempted.
- **Status** — Operation results.
Clicking details will also allow you to view errors, if any, that occurred during deployment.

Chapter 8 **Setting Deployment Options**

This chapter explains how to use the Configuration Builder to manage deployment options.

Topics

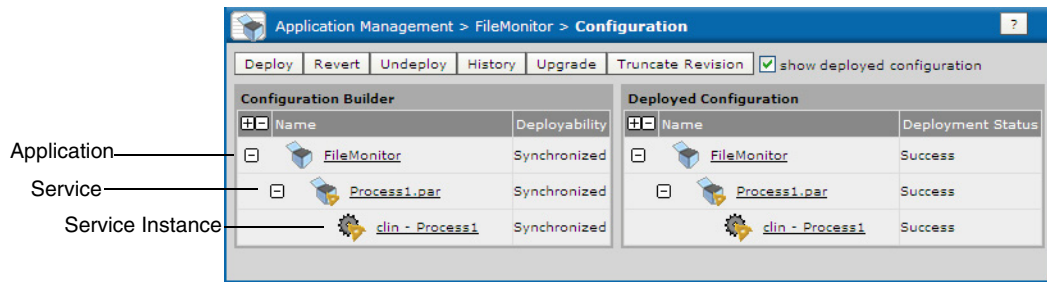
- [Configuration Console Overview, page 163](#)
- [Changing Global Variables at Deployment, page 164](#)
- [Setting Application Repository Instance Options, page 167](#)
- [Enabling a Process or Service to Run on Other Machines, page 168](#)
- [Adding a Custom Rulebase to a Process or Service, page 170](#)
- [Adding an Event to a Process or Service, page 174](#)
- [Setting Fault Tolerant Options for a Process, page 178](#)
- [Changing Global Variables for a Process or Service, page 184](#)
- [Configuring Storage for TIBCO BusinessWorks Processes, page 185](#)
- [Changing TIBCO BusinessWorks Process Configuration Properties, page 188](#)
- [Changing Server Settings, page 193](#)
- [Setting Graceful Shutdown Properties for a Process Engine, page 194](#)
- [Specifying HTTP Servlet Authentication Information, page 195](#)
- [Application Management Configuration Dialog, page 196](#)
- [Edit Application Configuration Dialog, page 199](#)
- [Edit Service Configuration Dialog, page 204](#)
- [Edit Service Instance Dialog, page 209](#)

Configuration Console Overview

When you create an application, the enterprise archive file you import has values defined for global variables. The process engines and service instances in the archive have configuration options set as well. When you deploy the application, you can use the options set in the archive, or change options in TIBCO Administrator.

The Configuration console consists of two panes, Configuration Builder and Deployed Configuration. Each pane contains applications, service configurations, and service instances as shown in the next diagram.

Figure 39 Configuration Console



The Configuration Builder pane at left allows you to deploy or update applications and to revert a deployment, that is, choose an earlier deployment configuration if there was one. You can deploy the same application multiple times, for example, to try out different machine configurations. However, only one deployment configuration can be running at any time. If you later wish to return to a previous deployment configuration, you can do so by choosing **Revert**.

You can also view the deployment history, the current deployment if there is one, and undeploy the application. See [Undeploying an Application on page 145](#) for a detailed discussion.

If you have installed new TIBCO software on a machine that is running process engines or service instances, you can upgrade them to use the new software by clicking **Upgrade**. See [Upgrading an Application on page 151](#) for more information.

When you select an application, service or service instance in the Configuration Builder panel, the displayed dialog allows you to change parameters for the deployment. When you select an application, service or service instance in the Deployed Configuration panel, the displayed dialog is read-only, providing a description of the properties.

Changing Global Variables at Deployment

An application's EAR file can contain global variables with values set at design-time. The global variables can be changed at deployment at the application level, service level, or service instance level.



Global variables are saved as XML data and must contain valid XML characters.

Global variable values can be set at the application, service, and service instance levels. A variable value set at the service instance level overrides the same variable value set at the service level. Similarly, a variable value set at the service level overrides the same variable set at the application level.

Each service instance can use the same variable and assign it a different value. The instance variable can be any services settable variable defined at configuration.



- If you change the default value of a global variable on deployment, that global variable is not changed by future deployments. That is, uploading a new EAR file does not change the values, even if the value of the global variable in the new EAR file is changed.
- If you do not change the default value of a global variable on deployment, the global variable's value is overwritten with the one in the EAR file when the new version of the EAR file is uploaded.
- Whenever the EAR file has the same value as the value in the TIBCO Administrator GUI, the value gets changed to whatever value is present in the EAR file on next deployment.

The rules for resolving global variable values are as follows.

- If the variable value is set at the service instance level in the TIBCO Administrator GUI and the value is different than the variable value set in the EAR file or the variable is not deployment settable, the value set at the service instance level is written to the deployed service instance's .tra file.

For example, for the application below, MYVARIABLE is set to 3 in the EAR file. MYVARIABLE is changed for each service instance in the TIBCO Administrator GUI to a different value. After deployment, the three service instances will use the value assigned in the TIBCO Administrator GUI for MYVARIABLE.

```
Application1 (MYVARIABLE 3)
  Service1 (MYVARIABLE 3)
    ServiceInstance1 (MYVARIABLE 4)
    ServiceInstance2 (MYVARIABLE 5)
    ServiceInstance3 (MYVARIABLE 6)
```

- If the value at the service instance level is the same as the value inside the EAR file, but the value set at the service level in the TIBCO Administrator GUI is different, then the value set at the service level is written to deployed service instance .tra file.

For example, for the application below, MYVARIABLE is set to 3 in the EAR file. MYVARIABLE is changed to 4 for Service1 and left unchanged for the ServiceInstance1 and ServiceInstance2. MYVARIABLE is changed to 5 for ServiceInstance3. After deployment, ServiceInstance1 and ServiceInstance2 will use 4 for MYVARIABLE. ServiceInstance3 will use 5 for MYVARIABLE.

```
Application1 (MYVARIABLE 3)
  Service1 (MYVARIABLE 4)
    ServiceInstance1 (MYVARIABLE 3)
    ServiceInstance2 (MYVARIABLE 3)
    ServiceInstance3 (MYVARIABLE 5)
```

- If the variable value set at the service instance level and the variable value set at the service level are the same as the value in the EAR file, but the value set at the application level is different, that variable is not written to the .tra file. At runtime, the application gets the value at the application level as this value is written to the application's repository.

For example, for the application below, MYVARIABLE is set to 3 in the EAR file. MYVARIABLE is changed to 4 for Application1 and is unchanged for all services and service instances. After deployment, all ServiceInstances will use 4 for MYVARIABLE.

```
Application1 (MYVARIABLE 4)
  Service1 (MYVARIABLE 3)
    ServiceInstance1 (MYVARIABLE 3)
    ServiceInstance2 (MYVARIABLE 3)
    ServiceInstance3 (MYVARIABLE 3)
  Service2 (MYVARIABLE 3)
    ServiceInstance11 (MYVARIABLE 3)
    ServiceInstance22 (MYVARIABLE 3)
    ServiceInstance33 (MYVARIABLE 3)
```

- If the value of the variable at all the three levels, application, service and service instance, is identical to the one inside the EAR file, the behavior is the same as that described in the previous bullet. However, if the variable is not deployment settable, the default value will be written to the .tra file.
- If there are no changes to the values of variables from the default values which are set in designer, the subsequent uploading of the EAR file will reset the variables default and configured values by new designer values of the new EAR file.

To Change Application Properties

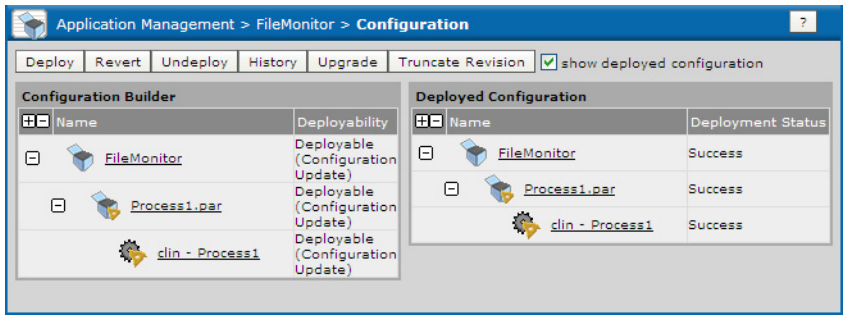
1. In the TIBCO Administrator GUI, click **Application Management**.
2. Click *Application* > **Configuration**, where *Application* is the application created when you loaded the enterprise archive file. See [Creating an Application on page 137](#).
3. In the Configuration Builder pane, select the level at which you want to change the variable, the *Application* name, service name or service instance name.

Click the **Advanced** tab to change values set for global variables. Use valid XML characters only. See [Global Variables on page 199](#) for descriptions.

The **Reset to Defaults** button restores all properties to the values defined in the enterprise archive file.

4. Click **Save**.
5. After you've made your changes, the Configuration Builder indicates that the deployment is out of date.

Figure 40 Changing Application Properties



6. Click **Deploy**.



If you have overridden the default value of a global variable, that global variable is not changed by future deployments, even if it is different than the variable value set in the new EAR file.

Setting Application Repository Instance Options

When TIBCO Administrator deploys an application, it creates an application repository instance that contains information about the application configuration. For administration domains that use TIBCO Rendezvous as the transport, the instance can be managed by the administration server or sent to the target machine on which the client application is deployed. Domains that use TIBCO Enterprise Message Service as the transport always send the application repository instance to the target machine.

- If the application repository instance is managed by the administration server, you can set the transport used by the client application to communicate with the administration server. The choices are rv (TIBCO Rendezvous), or HTTP (or HTTPS if the administration server is configured for it).
- If your domain was configured to send the application repository to the target machines where applications are run, the default choice is local. If this choice is used, each target machine must have TIBCO Runtime Agent 5.3 (or later) installed.

See the *TIBCO Administrator Server Configuration Guide* for more information.

To Change Application Repository Instance Options

You can only change application repository instances options for domains that use TIBCO Rendezvous.

1. In TIBCO Administrator, click **Application Management**.
2. Click *Application* > **Configuration**, where *Application* is the application created when you loaded the enterprise archive file. See [Creating an Application on page 137](#).
3. In the Configuration Builder pane, select the *Application* name.
4. Click the **Advanced** tab to change values for the repository instance. See [Edit Application Configuration Dialog on page 199](#) for more information.
5. Click **Save**.

Enabling a Process or Service to Run on Other Machines

You can assign a process or service to run on any machine that is part of your administration domain. See the *TIBCO Runtime Agent Domain Utility User's Guide* for information about adding a machine to a domain.



Adding a process to additional machines is useful for fault tolerance. As a rule, it therefore does not make sense to run the same process on the same machine twice.

TIBCO Adapter service instances cannot be assigned fault tolerant options.

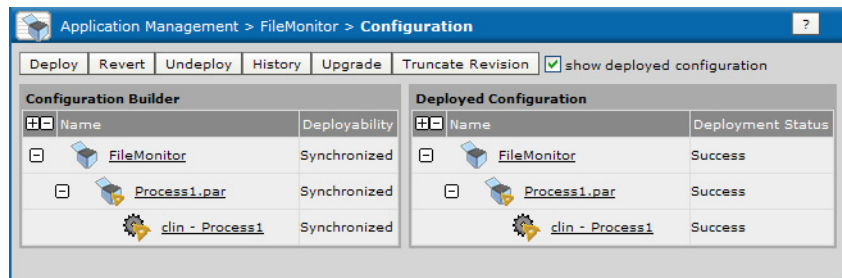
A service can be enabled or disabled. Only enabled services are deployed. When you disable a service, it is no longer deployed the next time you deploy the application, while all other services in the application are deployed as before. This can be useful, for example when you wish to deploy an application that includes a service for which you don't have the required software.

Only machines that have the software required by the process or service are visible when selecting the machine.

To Enable a Process or Service to Run on Other Machines

1. In TIBCO Administrator, click **Application Management**.
2. Select an application and expand it.

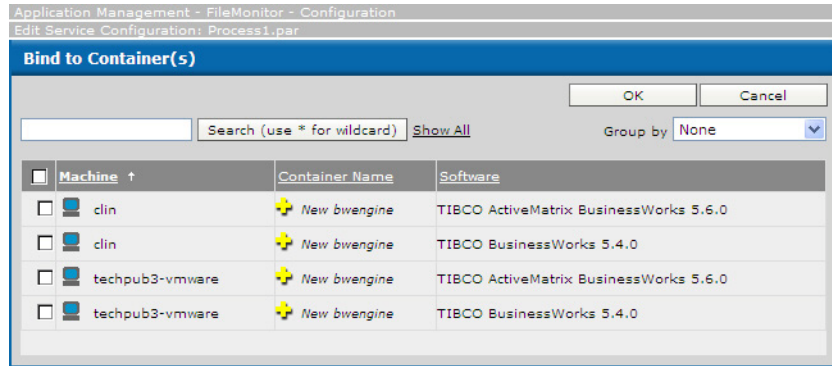
Figure 41 Enabling a Process or Service to Run on Other Machines



3. In the Configuration Builder pane, click a service or process name. A service is named with a .aar suffix. A process is named with a .par suffix.
4. In the General pane enable or disable the process or service by selecting or clearing the Enable Service check box.
5. In the Target Machines pane, click **Add to Additional Machines** to add a selected process or service to another machine.

- A dialog appears, similar to the following, displaying all machines in the domain on which the software required by the process or service is available. Select one or more machine, and then click **OK**.

Figure 42 *Select One or More Machine*



- Click **Save**.

See Also

See [Edit Service Configuration Dialog on page 204](#) for more information.

See [Setting Fault Tolerant Options for a Process on page 178](#).

See [Configuring Fault-Tolerant Engines on page 179](#).

Adding a Custom Rulebase to a Process or Service

The TIBCO Hawk agent monitors managed objects by processing rulebases, which are named collections of rules that contain management logic. Using TIBCO Hawk Display, you can create rulebases with specialized rules. (TIBCO Hawk Display is not included in TIBCO Runtime Agent). Hawk allows you to specify a very large number of alert conditions and alert results. You must have purchased the full TIBCO Hawk product to create TIBCO Hawk rulebases.

The same rulebase can be loaded on a single service, or multiple services.

Multiple rules defined in the same rulebase can monitor a particular application or system function. For example, an application rulebase could include one rule for issuing a medium-level alert if disk space or CPU usage exceeds certain thresholds. Another rule could issue a high-level alert and send a pager message to the system administrator if the application process terminates.

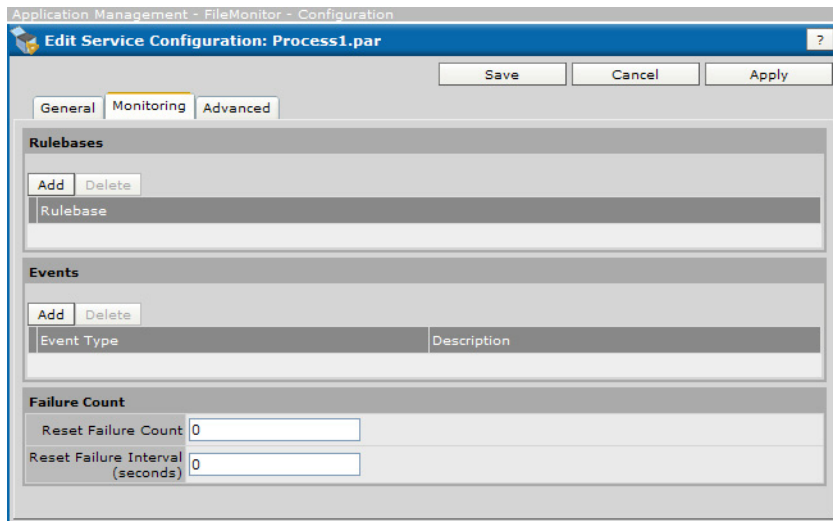
To Add a Custom Rulebase to an Application

This section provides information about adding a rulebase for a service or process. Information about building the rulebase expression is not provided. See the *TIBCO Hawk Administrator's Guide* for information about creating rulebases. The guide is part of the TIBCO Hawk documentation set.

1. In TIBCO Administrator, click **Application Management**.
2. Select the application for which the rulebase has been defined, and expand it.
3. In the Configuration Builder pane, click the service or process name for which the rulebase has been defined. A service is named with a `.arr` suffix. A process is named with a `.par` suffix.
4. Click the **Monitoring** tab.
5. In the Rulebases panel, click **Add**.
6. Click **Browse** and in the window that appears, navigate to the directory where the rulebase is stored and select the rulebase. Click **OK**.
7. Click **Save**.

For example, the next diagram shows the rulebase section for a process archive.

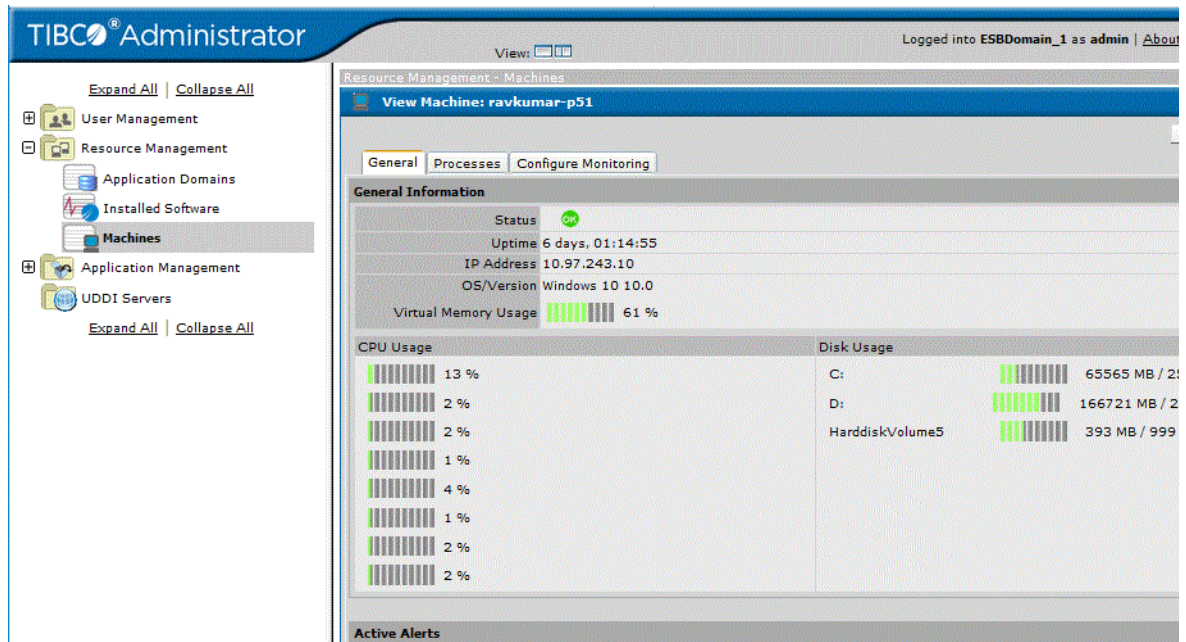
Figure 43 Adding a Custom Rulebase to an Application



When you deploy the service, TIBCO Hawk agent saves the rulebase file in the `TIBCO_HOME/domain/domain/rulebase` folder. The original rulebase can be safely removed, because the rulebase has been loaded into the domain. If you change the original rulebase, it must be reloaded into the service and the application must be redeployed.

When the conditions specified in the rulebase occur, the results display in the Resource Management > Machines View Machine panel. For example, the next screen shows several alerts that have been generated based on a rulebase.

Figure 44 Results Displayed When Conditions Specified in the Rulebase Occur



How to Create a Custom Rulebase

You can create rulebases using the TIBCO Hawk Display rulebase editor. The rulebase file name can be saved without using a naming convention (as was necessary in releases prior to 5.3). TIBCO Hawk agent creates the appropriate rulebase name and file when the service instance to which the rulebase is assigned is deployed.



The rulebase file name should not contain the space character.

For example, if two rulebase files are created and named:

- shared_custom1.hrb
- shared_custom2.hrb

And the above rulebase files are assigned to two service instances (as described in the previous section).

- D1-Process_Archive

- E1-Process_Archive

When the D1-Process_Archive service is deployed, TIBCO Hawk agent creates the following rulebase files for the service.

- D1-Process_Archive-shared_custom1.hrb
- D1-Process_Archive-shared_custom2.hrb.

Similarly, when the D1-Process_Archive-1 service is deployed, TIBCO Hawk agent creates the following rulebase files for the service:

- E1-Process_Archive-shared_custom1.hrb
- E1-Process_Archive-shared_custom2.hrb.

The rulebase file names for each instance are stored in an external property file so the TIBCO Hawk agent knows where to re-load the rulebase files if it is restarted.

- Rulebase file names are stored in the `TIBCO_HOME/tra/domain/domain/startup/application.properties` file in the `rbList` property.
- TIBCO Hawk agent assumes that the input rulebase file name provided when uploading a rulebase file uses the `.hrb` extension. If there is no file extension, Hawk Agent appends `.hrb` to the rulebase file during deployment.

Variable Substitution

You can assign certain variables to a rulebase and TIBCO Hawk agent will substitute values for the variables when the application is deployed. Variable substitution is typically used in a rulebase to change the data source from pointing to a specific service instance to point a generic service instance.

The following variables are supported by TIBCO Hawk agent:

- `%%TIBCO_DEPLOYMENT%%`— When encountered, the Hawk Agent substitutes the application's deployment name.
- `%%TIBCO_COMPONENT_INSTANCE%%` — When encountered, Hawk Agent substitutes the service instance name.
- `%%TIBCO_DOMAIN%%`— When encountered, Hawk Agent substitutes the administration domain name.
- `%%TIBCO_COMPONENT_TYPE%%`— When encountered, Hawk Agent substitutes the component's type.

Adding an Event to a Process or Service

An event can be configured without installing the TIBCO Hawk product. You can create events that generate an alert, cause email to be sent, or execute a command.

To Add an Event to a Service

You can define an event type to respond to a service instance failure, or to be triggered when a match occurs for some condition that is reported in the service instance log file.

1. In TIBCO Administrator, click **Application Management**.
2. Select an application and expand it.
3. In the Configuration Builder pane, click a service or process name. A service is named with a `.arr` suffix. A process is named with a `.par` suffix.
4. Click the **Monitoring** tab.
5. Click **Add** in the Events panel.
6. Specify the conditions and the event.
 - a. First choose a condition in the General panel.
 - b. In case the condition is met, you can choose to send an alert, send an email, or execute a command, or a combination of those.
7. When you've configured both condition and event, click **OK**.
8. Click **Save**.

For example, the next diagram shows the Add Event panel for a process archive.

Figure 45 Add Event Panel for a Process Archive

Application Management - File Adapter - AdFiles_delimitedReader - Configuration
 Edit Service Configuration: delimitedReader.aar

Add Event

OK Cancel

General

Event Type: Any Failure

Restart Service Instance: ☐

Description:

Alert

Generate Alert: ☐

Generate Alerts For: ☒ First Occurrence ☐ All Occurrences

Level: Low

Message:

Email

Send Email: ☐

Send Email For: ☒ First Occurrence ☐ All Occurrences

Mail Server:

To:

CC:

Subject:

Message:

Command

Execute Command: ☐

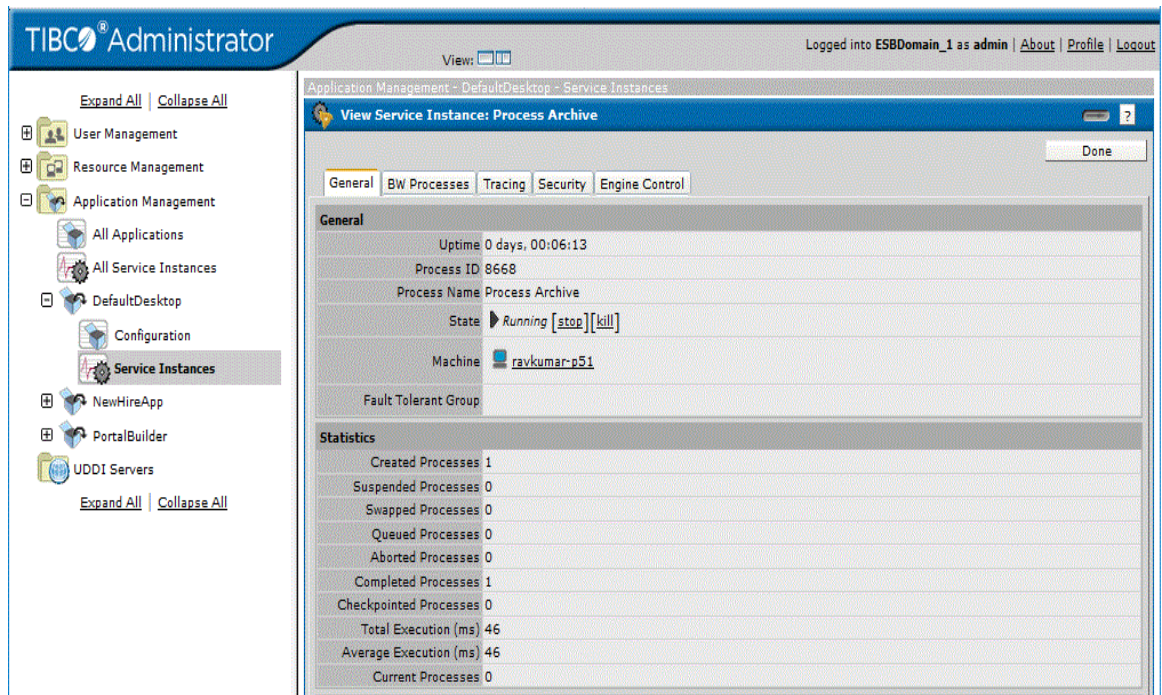
Execute Command For: ☒ First Occurrence ☐ All Occurrences

Command:

Arguments:

The event can be sent as an alert, by email or can trigger an operating system command. If the event is sent as an alert, it appears in the View Service Instance dialog under the Active Alerts pane. For example:

Figure 46 View Services Instances



See Also

See [Edit Service Configuration Dialog on page 204](#) for more information.

See [Specifying a Custom Alert on page 264](#) for an example of event configuration.

To set Events for a Process

If you are setting events for a process, the following fields display:

Any Failure

The event will be executed each time on any failure.

- Restart Service Instance restarts this service instance.
- Description is the description used in the event (email, alert, etc.).

First Component Failure

The event will be executed the first time the service fails.

- Restart Service Instance restarts this service instance.
- Description is the description used in the event (email, alert, etc.).

Second Component Failure

The event will be executed the second time the service fails.

- Restart Service Instance restarts this service instance.
- Description is the description used in the event (email, alert, etc.).

Subsequent Component Failure

The event will be executed all subsequent times the service fails.

- Restart Service Instance restarts this service instance.
- Description. The description used in the event (email, alert, etc.).

Suspended Process

The event will be executed when the service is suspended. For example, a process can be suspended by a Wait activity. Note that alerts associated with the suspended process will be cleared after 15 minutes.

- Restart Service Instance restarts this service instance.
- Description is the description used in the event (email, alert, etc.).



The event type "Suspended Process" does not work for Adapter Services.

Log event

The event will be executed each time a log event matching the Match field occurs.

- Restart Service Instance restarts this service instance.
- Match is the string TIBCO Administrator looks for in the log event.
- Description is the description used in the event (email, alert, etc.).

Setting Fault Tolerant Options for a Process

The FT Group Settings panel displays only if the TIBCO BusinessWorks process you have selected has been added to at least two (different) machines. If your domain includes components that were deployed as part of a fault-tolerant group, the display includes the information about the group.



Fault tolerance options can be set only for TIBCO BusinessWorks processes. TIBCO Adapter services cannot be assigned fault tolerant options.

You can start one or more process engines in the group. If more than one engine has started, only one is displayed as Running and all other engines are displayed as Standing By (or, initially, as Starting Up).

When you change the status of a component that has been deployed as part of a FT group, the status change affects all other members of the group.

- After you have deployed the process engines, it is most efficient to select all process engines by selecting the check boxes, and then choosing **Start**. After the primary and secondary engines have communicated, the master will display as Running and all other engines as Standby. If you start only the primary, it will first go to Standby mode as it checks the status of the other engines. It then changes to Running.
- If you shut down a process engine, the appropriate secondary engine starts automatically.

To Set Fault Tolerant Options

1. In TIBCO Administrator, click **Application Management**.
2. Select an application and expand it.
3. In the Configuration Builder pane, click process name. A process is named with a `.par` suffix.
4. Click the **General** tab.
5. Select **Run Fault Tolerant**. Change other options as required. See [FT Group Settings on page 205](#) for field descriptions.
6. Click **Save**.

Changing the Checkpoint Data Repository for a Process

To run TIBCO BusinessWorks using multiple engines in fault tolerant mode, you must specify a checkpoint data repository. See [Failover and Checkpoint Data on page 182](#) for more information.

For true fault tolerance, you must store the data in a database. You specify a JDBC Connection resource for the database to be used when you configure your project in TIBCO Designer. The database is then one of the available options on the Checkpoint Data Repository pop-up menu.

To Change Checkpoint Data Repository Properties

1. In TIBCO Administrator, click **Application Management**.
2. Select an application and expand it.
3. In the Configuration Builder pane, click a process name. A process is named with a .par suffix.
4. Click the **Advanced** tab.
5. Change properties as required.
6. Click **Save**.

Configuring Fault-Tolerant Engines

The TIBCO BusinessWorks process engine can be configured to be fault-tolerant. You can start up several engines. In the event of a failure, other engines restart process starters and the corresponding services.

If you use a database to store process engine information, a service is reinstantiated to the state of its last checkpoint. In the event of a failure, any processing done after a checkpoint is lost when the process instance is restarted by another engine. See *TIBCO BusinessWorks Palette Reference* for more information about Checkpoint activities. See [Configuring Storage for TIBCO BusinessWorks Processes on page 185](#) for more information about configuring process engine storage.

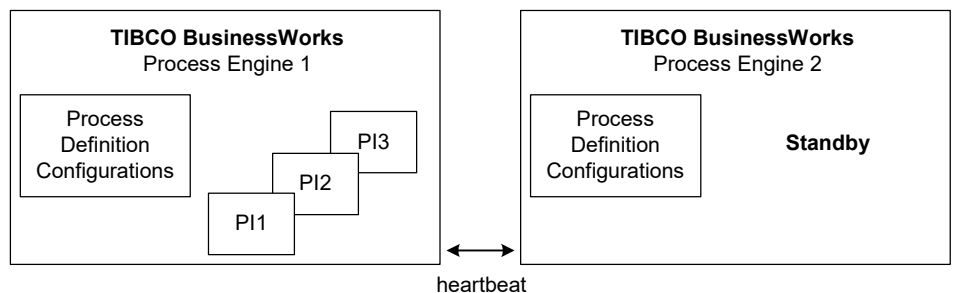


When applications are deployed using

- local transport, BW engine fault tolerance does not rely on the administrator server. The BW engines run without accessing data from the administrator server.
- rv or http transport options, the administration server must be running for BW fault tolerance to work correctly. If the master BW engine fails, a running and fully initialized secondary BW engine will take over. However, if all instances of the administrator server are down, BW engines cannot be restarted (they cannot get repository data from the server).

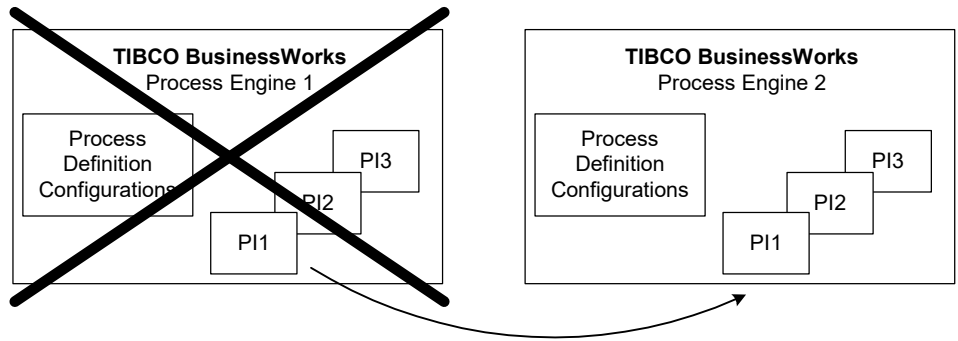
The next diagram illustrates normal operation of a fault-tolerant configuration. One engine is configured as the master, and it creates and executes services. The second engine is a secondary engine, and it stands by in case of failure of the master. The engines send heartbeats to notify each other they are operating normally.

Figure 47 Normal operation: master processing while secondary stands by



In the event the master process engine fails, the secondary engine detects the stop in the master's heartbeat and resumes operation in place of the master. All process starters are restarted on the secondary, and services are restarted to the state of their last checkpoint. The next diagram illustrates a failure and the secondary restarting the service

Figure 48 Fault-tolerant failover



The expected deployment is for master and secondary engines to reside on separate machines. You can have multiple secondary engines, if desired, and you can specify a weight for each engine. The weight determines the type of relationship between the fault-tolerant engines. See [Peer or Master and Secondary Relationships on page 181](#) for more information about relationships between fault-tolerant engines.

A master and its secondary engines is known as a *fault-tolerant group*. The group can be configured with several advanced configuration options, such as the heartbeat interval and the weight of each group member. See *TIBCO BusinessWorks Palette Reference* for a complete description of configuration options for fault tolerance.

Peer or Master and Secondary Relationships

Members of a fault-tolerant group can be configured as peers or as master and secondary engines. If all engines are peers, when the machine containing the currently active process engine fails, another peer process engine resumes processing for the first engine, and continues processing until its machine fails.

If the engines are configured as master and secondary, the secondary engine resumes processing when the master fails. The secondary engine continues processing until the master recovers. Once the master recovers, the secondary engine shuts down and the master takes over processing again.

The Fault Tolerance tab of the Process Engine deployment resource allows you to specify the member weight of each member of a fault-tolerant group. The member with the highest weight is the master. You can select "Peer" in the first field on the tab to configure all engines as peers (that is, they all have the same weight). You can select Primary/Secondary to configure the engines as master and secondary. You can also select Custom to specify your own values for the weight of each member of the group.

Failover and Checkpoint Data

A checkpoint saves the current state of a running process instance. For a secondary process engine to resume running process instances from their last checkpoint, the secondary process engine must have access to the saved state of the process instances from the master process engine.

If you select the service (.par), and then the Advanced tab, a pane named TIBCO BusinessWorks Checkpoint Data Repository is displayed. In this field, you can specify where state of process instances is stored when a checkpoint is performed. The value defaults to Checkpoint Data Repository. If a JDBC Connection Resource has been configured for the project, you also have the option to choose database.

Because fault-tolerant engines are expected to be on separate machines, you should specify to use a database for storage for each process engine. By this you can specify the same JDBC Connection resource for the master and secondary engines, and therefore all engines can share the information stored for process instance checkpoints. See [, Configuring Storage for TIBCO BusinessWorks Processes, on page 185](#).

If all engines share the checkpoint information, and then the secondary engines can recover process instances up to their last checkpoint. If engines do not share the checkpoint information, process instances are not restarted.

Process Starters and Fault-Tolerance

When a master process engine fails, its process starters are restarted on the secondary engine. This may not be possible with all process starters. For example, the HTTP Receiver process starter listens for HTTP requests on a specified port on the machine where the process engine resides. If a secondary engine resumes operation for a master engine, the new machine is now listening for HTTP requests on the specified port. HTTP requests always specify the machine name, so incoming HTTP requests will not automatically be redirected to the new machine.

Each process starter has different configuration requirements, and not all process starters may gracefully resume on a different machine. You may have to provide additional hardware or software to redirect the incoming events to the appropriate place in the event of a failure.

Also, your servers may not have all of the necessary software for restarting all of instances. For example, your database may reside on the same machine as your master process engine. If that server goes down, any JDBC activities will not be able to execute. Therefore, you may not wish to load process definitions that use JDBC activities in your secondary process engine.

You can specify that your secondary process engine loads different process definitions than the master. You may only want to load the process definitions that can gracefully migrate to a new server during a failure.

See Also

See [FT Group Settings on page 205](#) for field descriptions.

For information about which process engine starts and what state it starts in, see *TIBCO BusinessWorks Process Design Guide*.

Changing Global Variables for a Process or Service

Some service and process specific global variables can be defined in TIBCO Designer, and changed in TIBCO Administrator. When defining global variables in TIBCO Designer, you specify whether the variable should be settable at design-time only, for the deployment, or for the service. When configuring a service or process, all variables that were designated settable for the service are then displayed in TIBCO Administrator.



Global variables are saved as XML data and must contain valid XML characters only.

Click **Reset to Defaults** to reset to the default values defined in the enterprise archive file.

To Change Global Variables for a Service or Process

1. In TIBCO Administrator, click **Application Management**.
2. Select an application and expand it.
3. In the Configuration Builder pane, click a service or process name. A service is named with a `.arr` suffix. A process is named with a `.par` suffix.
4. Click the **Advanced** tab.
5. Change global variables as required. Be sure to use valid XML characters only.
6. Click **Save**.

See Also

Global variables are described in your TIBCO Adapter and TIBCO BusinessWorks documentation sets.

Configuring Storage for TIBCO BusinessWorks Processes

You can use TIBCO Administrator to configure the location where TIBCO BusinessWorks process engines store internal information.

Most of the information a process engine stores is information about each service's state when a checkpoint is taken. There is, however, some other internal information stored by the engine. You can specify that this information is stored in the file system (the default) or in a database.

For some systems, using a file system for storage may be sufficient. However, some functionality is only available when you use a database for storing information about service state:

- With a database for storage, fault-tolerant engines can recover services up to a checkpoint. Without a database, running services cannot be recovered to their last checkpoint.
- With a database for storage, Wait/Notify activities can be used to pass data between services running on different machines. Without a database, the wait and notify activities cannot communicate across machines.

Specifying a Database for Storage

To configure a database for storage, follow these steps:

1. In TIBCO Designer, be sure to specify a JDBC Connection resource for the database you wish to use, and then build the EAR file.
2. After you have uploaded the EAR file and created the application in the TIBCO Administrator GUI, select Application Management then select the application in the Configuration Builder pane of the Configuration console.
3. Select the service (.par) in the Configuration Builder pane and choose the Advanced tab.

Figure 49 Specifying a Database for Storage



4. In the TIBCO BusinessWorks Checkpoint Data Repository pane, select the Database pane.

Database Table Names

When you specify a database for TIBCO BusinessWorks storage, tables are created in your database. The administration domain name and deployment ID (assigned by TIBCO BusinessWorks) are used to name the tables to ensure uniqueness of the tables for each domain and each deployment.

Because some databases limit the number and kinds of characters for table names, the domain name can altered before being used in the table name. The first eight characters and the last eight characters of the domain name are taken and any non-legal characters (such as spaces or dashes) are converted to underscores. This creates a sixteen-character unique ID for each domain, provided that the first and last eight characters of all of your domain names are unique.

For example, the following illustrates conversion of domain names. Notice the second and third domain names convert to the same ID. You should avoid this by creating domain names so that the combination of the first and last eight characters are unique.

Table 5 Conversion of Domain Names

Domain Name	Converts To Domain ID
TIBCO_domain_Accounting	TIBCO_docounting
TIBCO_domain_Marketing	TIBCO_doarketing
TIBCO_domain_Direct_Marketing	TIBCO_doarketing

All table names created by TIBCO BusinessWorks begin with *APPLICATION_NAME_random number*. You can alter the storage parameters for these tables if you desire, but the table names and column definitions must remain the same.

Manually Creating Database Tables

The process engine creates database tables used to store process engine information automatically. Some database administrators do not permit applications to automatically create tables. If you wish to manually create the database tables, TIBCO BusinessWorks provides template scripts for the supported databases in the `TIBCO_BW_HOME/bin` directory.

In these scripts `TABLE_NAME_PREFIX` and `ENGINE_NAME_MAX_LENGTH` are placed in the SQL code as placeholders. The `TABLE_NAME_PREFIX` is determined by default at deployment time (see [Database Table Names on page 186](#) for a description of how the table name prefix is determined). You can obtain this prefix by locating

the `Database.Tablename.Prefix` property in the deployment configuration file and substituting its value where required in the template SQL script. `ENGINE_NAME_MAX_LENGTH` is 128, so supply that value instead of the placeholder in the SQL script.

To Manually Create the Tables

1. Create a deployment configuration that specifies a database for process engine storage. See [Specifying a Database for Storage on page 185](#).
2. Before starting the process engine, examine the deployment configuration file (the *processEngine.tra* file) and locate the property `Database.Tablename.Prefix`.
3. Edit the appropriate SQL script template for the database you are using and replace `TABLE_NAME_PREFIX` with the value of the `Database.Tablename.Prefix` property.
4. Change `ENGINE_NAME_MAX_LENGTH` to 128.
5. Save the changes to the SQL script.
6. Run the SQL script against the database you wish to use.
7. Start the process engine.

Changing TIBCO BusinessWorks Process Configuration Properties

You can enable or disable TIBCO BusinessWorks processes, set the maximum jobs that can be in memory, enable or disable the activation limit and set the maximum jobs that can be in memory and paged.

To Change Process Configuration Properties

1. In TIBCO Administrator, click **Application Management**.
2. Select an application and expand it.
3. Click **Configuration**.
4. In the Configuration Builder pane, click a process name. A process is named with a .par suffix.
5. Click the **Advanced** tab.
6. Change properties as required. The remaining topics in this section provide information about the properties you can set.
7. Click **Save**.

Controlling Execution of TIBCO BusinessWorks Services

Process starters create process instances to handle incoming events. Process instances consume memory and CPU resources on your system. Depending on the available machine resources, you may only be able to run a limited number of process instances concurrently.

Process instances typically remain in memory as long as they are executing an activity. If the process instance is waiting for an incoming event (for example, a Wait for Adapter Message activity), the process instance can be paged out to disk and resumed later after the event arrives. New process instances are paged out to disk until there is available memory and resources to accommodate them.

You can use TIBCO Administrator to control the execution of TIBCO BusinessWorks process instances. This is useful if your system has limited memory or resources, or if you want to restrict process instances to run sequentially.

The TIBCO BusinessWorks Process Configurations dialog, accessed by selecting the process archive in the Configuration Builder, allows you to specify the following:

- **Max Jobs** — Specifies the maximum number of process instances that can concurrently be loaded into memory.
- **Activation Limit** — Specifies that once a process instance is loaded, it must remain in memory until it completes.
- **Max Jobs (Flow Limit)** — Specifies the maximum number of currently running process instance to start before suspending the process starter.

The following sections describe these configuration settings and the relationship between the settings.

To Access the Process Configurations

1. Select the process archive (.par) in the application's Configuration Builder pane.
2. Click the Advanced tab and the + next to the process archive's name.

Specifying the Maximum Number of Concurrently Active Processes

Incoming events may not be evenly distributed over time. That is, there may be periods where a large number of incoming events occur and other periods where relatively few events occur. To prevent your system from being overwhelmed by incoming events, the Max Jobs (Flow Limit) field limits the number of process instances created by a process starter. This allows you to control the flow of processing so that incoming events are no longer accepted when the limit is reached.

Controlling the flow of processing is especially useful when you are using protocols that can store unsent messages on the server until the receiver is ready to process them. For example, if your process definition polls an email server for new messages (that is, Receive Mail is the process starter), and then you can set Max Jobs (Flow Limit) to control the number of process instances created for each new email. Email that has not been processed remains on the email server until the process engine is ready to create more process instances. Other protocols where this approach are useful are TIBCO Rendezvous Certified Messaging (RVCM), JMS durable topic subscriptions, and JMS queues.

When a process engine reaches the specified Max Jobs (Flow Limit), it is placed in a FLOW_CONTROLLED state. In this state, the process engine can continue executing existing process instances, but new process instances are not allowed. Incoming messages can then be directed to another process engine. A process engine will resume creating new process instances once a sufficient number of its current process instances have completed.

Specifying Maximum Number of Concurrent Processes in Memory

The Max Jobs field in the Process Configurations dialog allows you to specify the maximum number of concurrent process instances that can be stored in memory. For example, if you set Max Jobs to 5, the process engine can only keep five process instances in memory. Any process instances created once the maximum is reached must be paged out to disk.

Specifying a value for Max Jobs causes the process engine to incur some overhead for managing the paging of process instances to and from disk. If you have sufficient system resources and do not expect incoming events to exceed the limits of your system, consider specifying Max Jobs as 0. This allows the process engine to create an unbounded number of process instances and eliminates the overhead of paging.

Keeping Services in Memory

The Activation Limit field specifies that once a process instance is loaded into memory, it should not be paged out to disk until it completes. This option is useful if you wish to specify sequential processing of incoming events, or if you want to enforce limited concurrent execution of process instances.

Effects of Setting the Configuration Fields

The Max Jobs and Activation Limit options work together to provide different concurrency limits. The Max Jobs (Flow Limit) field also affects the concurrency limit. The next table describes the effects of various combinations of these options.

Table 6 *Effects of various configuration settings*

Max Jobs	Activation Limit	Max Jobs (Flow Limit)	Description
0	Cleared or selected	0	An unlimited number of process instances can be created and concurrently loaded into memory. Activation Limit is ignored when Max Jobs is set to 0.
0	Cleared or selected	N	No paging of process instances. Allows up to N process instances before placing process starter in flow controlled stated. Activation Limit is ignored when Max Jobs is set to 0.

Table 6 Effects of various configuration settings

Max Jobs	Activation Limit	Max Jobs (Flow Limit)	Description
1	Selected	N	<p>One process instance is loaded into memory at a time and kept there until it completes its execution. This guarantees incoming events are processed in the order in which they occur. Up to N process instances are paged to disk, and then the process starter is placed into flow controlled state.</p> <p>Note: If your goal is to sequentially process incoming events, use the Sequencing Key field on the Misc tab of the process starter. Using Max Jobs and Activation Limit incurs overhead as process instances are paged to disk and retrieved from disk.</p>
1	Selected	0	<p>Once process instance is loaded into memory at a time and kept there until it completes its execution. This guarantees incoming events are processed in the order in which they occur. There is no limit on the number of process instances that can be created and paged to disk.</p> <p>Note: If your goal is to sequentially process incoming events, use the Sequencing Key field on the Misc tab of the process starter. Using Max Jobs and Activation Limit incurs overhead as process instances are paged to disk and retrieved from disk.</p>
1	Cleared	N	<p>One process instance is loaded into memory at a time, but up to N process instances are created. Incoming events can be processed in any order because process instances are not kept in memory until they complete execution.</p>

Table 6 Effects of various configuration settings

Max Jobs	Activation Limit	Max Jobs (Flow Limit)	Description
M	Selected	0	<p>An unlimited number of process instances can be created, but only M are loaded into memory and processed concurrently.</p> <p>This setting ensures a limited amount of concurrent processing. This situation is useful if you have limited resources, such as database connections. You can set Max Jobs to a relatively small number and the Activation Limit option keeps each service in memory until the service completes. Each loaded process uses a machine resource until the service completes. Once a service releases the resource, a new process can be loaded into memory and the corresponding service can use the resource.</p>
		N	Same as above, except only N process instances are created before the process engine is placed in the flow controlled state.
M	Cleared	0	<p>An unlimited number of process instances can be created, but only M are loaded into memory and processed concurrently. After M process instances are created, new process instances are paged to disk. There is no guarantee of the order in which process instances are executed.</p>
		N	Same as above, except only N process instances are created before the process engine is placed in the flow controlled state.

Changing Server Settings

You can change the following properties for a process engine or service instance. You can also modify Java properties, such as changing the classpath and managing the heap size. In addition, you can set whether the instance should run as a Windows Service and define startup options.

- Start on Boot
- Enable Verbose Tracing
- Max Log File Size
- Max Log File Count
- Thread Count

To Change Server Settings for a Process or Service Instance

1. In TIBCO Administrator, click **Application Management**.
2. Select an application and expand Configuration.
3. Click the Process Archive.
4. Click the **Server Settings** tab.
5. Change options as required.
6. Click **Save**.

See Also

See [Server Settings Tab on page 209](#) for field descriptions.

Setting Graceful Shutdown Properties for a Process Engine

The graceful shutdown command causes the process engine to deactivate all process starters after you click the Stop button for an application, and wait (up to the maximum timeout) for all current jobs to either finish or take a checkpoint, before shutting down the engine. If Wait For Checkpoints is selected, the engine will wait up to the Kill Jobs Timeout for all jobs to finish even if they take a checkpoint.

To Set Graceful Shutdown Properties

1. In the TIBCO Administrator GUI, click **Application Management**.
2. Select an application and expand it to view a process engine.
3. Click a process engine.
4. Click the **Graceful Shutdown** tab.
5. Change options as required.
6. Click **Save**.

See Also

[Graceful Shutdown Tab on page 211.](#)

Specifying HTTP Servlet Authentication Information

Some products, such as TIBCO BusinessWorks Workflow, allow users to specify HTTP servlet authentication information in an HTTP servlet authentication pane that is part of the Advanced tab for a service or process.



This option is only available if your product supports this functionality. The available options also vary from product to product.

The HTTP Servlet Authentication pane allows you to choose an authentication type. The options you can select depend on the type you choose.

- **HTTP Session**—Allows you to specify an idle timeout for that session.

If you deploy a web application that is accessed by a web browser, the value defined in the HTTP Session field for the first web application is used. If multiple web applications are deployed, each subsequent web application will use the timeout value that was set for the first deployed web application. The timeout value set for each subsequent web application will be ignored.

- **Cookie**—Allows you to specify the following information:

Idle Timeout—Determines when the session is terminated if idle.

Domain—Domain to use for the cookie if the domain cannot be determined from the request.

Path—URL path to apply this cookie for.

Cookie expire—Length of time to keep the cookie valid. Values range from 15 seconds to never.

Signature Password—Password used to protect the logged in identity from being changed in a client's cookie file.

- **Web Server**—Allows you to specify the following information:

Idle Timeout—Determines when the session is terminated if idle.

Require New Session for Verify—Allows users to specify that a new browser session is required for login.

- **External**—Allows you to specify the following options:

Idle Timeout—Determines when the session is terminated if idle.

Authentication Handler Class Name—Class name in servlet path or in the enterprise archive file.

Application Management Configuration Dialog

The application management configuration dialog displays the following panes, side by side, if Show deployed configuration is selected.

- [Configuration Builder Pane](#)
- [Deployed Configuration Pane](#)

Configuration Builder Pane

Deploy

Click **Deploy** to deploy the application. The deploy dialog appears. See [Deploy Configuration Dialog on page 159](#) for more information.

Revert

When you revert an application, you select a different configuration of the currently deployed configuration. You can then decide to deploy this deployment configuration. If you do, service instances are stopped, updated, and restarted. Any component that is removed from a machine as a result of the revert is undeployed from that machine. See [Reverting to a Previously Deployed Application on page 144](#) for more information.

Undeploy

Click **Undeploy** to undeploy the application. When you undeploy a deployed application, TIBCO Administrator stops all running services and removes them from the list of services that can be started. In effect, it completely removes all traces of the deployment (with the exception of the logs). See [Undeploying an Application on page 145](#) for more information.

History

Click **History** to view the deployment history for this application. See [Viewing Application Deployment History on page 149](#) for more information.

Upgrade

If you have installed new TIBCO software on a machine that is running process or service instances, you can upgrade the instances to use the new software by clicking **Upgrade**. See [Upgrading an Application on page 151](#) for more information.

Show deployed configuration

Select to display the Deployed Configuration dialog box where you can view detailed information about the components deployed in the application.

Truncate Revision

Purge old revisions from the revision histories of the selected applications. See [Purging Application Revisions on page 150](#) for details.

Configuration List

Each component and service in the application is listed along with one of the following descriptors in the Deployability column

- Deployable, (Remove) — On Component. The last uploaded enterprise archive file does not contain this component. The component and all service instances will be removed from the application on deploy.
On Service Instance — The service instance has been deleted. This will take effect on deployment.
- Deployable, (New) — The component or service instance has never been deployed successfully. If all service instances are removed and new ones added, the component will be in this state.
- Deployable (Archive Update) — The last uploaded enterprise archive file has changes related to this component. Changes will take effect on deployment.
- Deployable (Configuration Update) — The last uploaded enterprise archive file had deployment descriptors updated (typically global variables) that effect this component.
- Deployable (Configuration Changes) — Changes have been made to the service instance configuration and will take effect on deployment.
- Deployable (Last Deploy Failed) — The last deployment failed. History should have details. Likely problems are the TIBCO Hawk agent needs to be started on the target machine, or TIBCO Rendezvous communication or configuration parameters are not correct.

- Synchronized — The configuration is correct. There have been no changes since last successful deployment.
- Needs configuration — You must select a component or service instance and then each tab. Workflow in particular requires this for some automatic configuration to be done. Must be remedied or the component must be disabled before deployment can succeed.
- Need to deploy in a Service Container — There are no service instances specified for the component. You must either disable it or assign at least one machine to component to enable deployment.
- Need to bind to a Service — Not currently used.
- Deployable, services require deployment — The undeploy command was run. All services are configured correctly and are ready for deployment.
- Deployable, containers require deployment — The component had a service instance modified, added or removed. The change will take effect on deployment.
- Services require configuration — A component has a service instance that needs to be configured. Deployment cannot be done until this is remedied or the component is disabled.
- Containers require configuration — Not currently used.
- Disabled — The component is marked disabled and will not be deployed. If deployment is attempted, the component will be undeployed when deployment is done.
- Disabled, will remove existing configuration — The component for the deployed service instance was marked Disabled. When deployment is done, the service instance will be undeployed.

Deployed Configuration Pane

Displays deployed components for this application and their status. Click each component to view detailed information about the deployed component.

Edit Application Configuration Dialog

Fields can be edited if this dialog is invoked from the Configuration Builder pane. If invoked from the Deployed Configuration pane, the fields are read only.

The following tabs are available:

- [General Tab](#)
- [Advanced Tab](#)

General Tab

Application Archive

Provides information about the enterprise archive file including the package name, version, description, creation date and owner.

Upload New EAR File

Allows you to replace the current enterprise archive file with an updated version.

Application Parameters

Allows you to specify the application name, associated deployment name, a simple description, and contact name for the application. It also allows you to specify the maximum number of revisions to keep for this application in the revision history.

Advanced Tab

The **Reset to Defaults** button resets all global variables to default settings as set in the enterprise archive file.

Global Variables

Displays the global variables set in the enterprise archive file for this application.



Global variables are saved as XML data and must contain valid XML characters only.

The following global variables are predefined by default:

- **DirLedger** — Used by the system when defining the path name of the TIBCO Rendezvous certified messaging ledger file. The default is the root installation directory.
- **DirTrace** — Used by the system to partially create the path name for log file used by the adapter. The default is the root installation directory.
- **HawkEnabled** — Used by the system to indicate whether TIBCO Hawk is used to monitor the adapter. True indicates that a Hawk microagent is defined for the adapter. False indicates the microagent is not to be used. Default is False.
- **JmsProviderUrl** — A JMS provider URL tells applications where the JMS daemon is located. Setting this value mostly makes sense in early stages of a project, when only one JMS daemon is used.
- **JmsSslProviderUrl** — Specifies where the JMS server, running in the SSL mode, is located. Setting this value mostly makes sense in the early stages of a project, when only one JMS server is used.
- **RemoteRvDaemon** — Used by the system to identify the TIBCO Rendezvous routing daemon. See *TIBCO Rendezvous Administration* for details about specifying the routing daemon name.
- **RvDaemon** — Used by the system to identify the TIBCO Rendezvous daemon parameter. The parameter instructs the transport object about how and where to find the Rendezvous daemon and establish communication. The default value is 7500, which is the default value used by the Rendezvous daemon. See *TIBCO Rendezvous Concepts* for details about specifying the daemon parameter.
- **RvNetwork** — Used by the system to identify the TIBCO Rendezvous network parameter. Every network transport communicates with other transports over a single network interface. On computers with more than one network interface, the network parameter instructs the TIBCO Rendezvous daemon to use a particular network for all outbound messages from this transport. See *TIBCO Rendezvous Concepts* for details about specifying the network parameter.
- **RvService** — Used by the system to identify the TIBCO Rendezvous service parameter. The Rendezvous daemon divides the network into logical partitions. Each transport communicates on a single service; a transport can communicate only with other transports on the same service. See *TIBCO Rendezvous Concepts* for details about specifying the service parameter. Default is 7500

- **RvaHost** — Used by the system to identify the computer on which the TIBCO Rendezvous agent runs. See *TIBCO Rendezvous Administration* for details about specifying the rva parameters.
- **RvaPort** — Used by the system to identify the TIBCO Rendezvous agent TCP port where the agent listens for client connection requests. See *TIBCO Rendezvous Administration* for details about specifying the rva parameters. Default is to 7501.
- **TIBHawkDaemon** — Used by the system to identify the TIBCO Hawk daemon parameter. See the *TIBCO Hawk Installation and Configuration* manual for details about this parameter. Default is the value that was set during domain creation (7474 by default).
- **TIBHawkNetwork** — Used by the system to identify the TIBCO Hawk network parameter. See the *TIBCO Hawk Installation and Configuration* manual for details about this parameter. Default is an empty string.
- **TIBHawkService** — Used by the system to identify the TIBCO service parameter. See the *TIBCO Hawk Installation and Configuration* manual for details about this parameter. Default is 7474.
- **MessageEncoding** — The message encoding set for the application. The default value is ISO8859-1, which only supports English and other western European languages that belong to ISO Latin-1 character set. After the project is deployed in an administration domain, the messaging encoding set at design time is overridden by the domain's encoding property. All the TIBCO components working in the same domain must always use the same encoding for intercommunication. See *TIBCO Administrator Server Configuration Guide* for more information.

TIBCO BusinessWorks and Adapters Deployment Repository Instance

When TIBCO Administrator deploys an application, it creates an application repository which contains information about the application configuration. You can view and change certain aspects of the application repository.

In Transport you select the transport the administration server uses to communicate with the client application. You can change this setting only if your administration domain uses TIBCO Rendezvous as the transport. Administration domains that use TIBCO Enterprise Message Service server always use the local choice.

Choose local, rv (TIBCO Rendezvous) or HTTP, (or HTTPS if the administration domain has been set up to use HTTPS). The default is set when creating a domain.

- **local**. By default, the transport is set to local. This means that the application repository will be sent to the target machine. This allows the application to run independently of the administration server.

If you change the transport from local to another value, the application repository will not be pushed to the target machine, and the application will communicate with the administration server at runtime.

The local choice is supported only if the target machines have installed TIBCO Runtime Agent 5.3 or later.

When you choose local, you can set the message encoding for the application. See MessageEncoding above for more information.

For more information about these choices, see the *TIBCO Administrator Server Configuration Guide*.

- rv. If selected, the client application will use TIBCO Rendezvous to communicate with the administration server. The following fields become available:
 - Server Name — administration server name.
 - Instance Name — Service instance name, that is, the instance of the service running on a particular machine.
 - User Name — User authorized for this application repository. Defaults to the user currently logged into Administrator.
 - Password — User's password.
 - Timeout — Amount of time in seconds allowed for completing a task, such as retrieving information from the server. Defaults to 600 seconds.
 - Service, Network, Daemon — TIBCO Rendezvous connection parameters used.
 - Discovery Timeout — Amount of time in seconds allowed for the initial connection to the administration server.
 - Regional Subject — TIBCO Rendezvous subject prefix used for regional read-operation in the load balancing mode. For more information see the *TIBCO Administrator Server Configuration Guide*.
 - Operation Retry — Number of times to retry after a communication timeout occurs.
- http. If selected, the client application will use HTTP to communicate with the administration server.

If your administration domain is not initially enabled for HTTPS, and there are deployed applications in the domain that use HTTP to connect to the application repository, the service instances will not restart after they are shut

down. In this case, you must redeploy each service instance after changing the transport to HTTPS.

- Server Name — administration server name.
- Instance Name — Service instance name, that is, the instance of the service running on a particular machine.
- User Name — User authorized for this application repository. Defaults to the user currently logged into Administrator.
- Password — User's password.
- Timeout — Amount of time in seconds allowed for completing a task, such as retrieving information from the server. Defaults to 600 seconds.
- HTTP URL, HTTPS URL — The URL on which the client attempts to connect to the server. What displays depends on whether you configured the server for HTTPS.



You cannot use HTTP or HTTPS to connect to a 4.x adapter.

Preview URL

If you have selected, rv or http in the Transport field, click the preview URL to display the URL that the application uses to access the application repository.

Edit Service Configuration Dialog

Fields can be edited if this dialog is invoked from the Configuration Builder pane. If invoked from the Deployed Configuration pane, the fields are read only.

The following tabs are available:

- [General Tab](#)
- [Monitoring Tab](#)
- [Advanced Tab](#)

General Tab

General

- Name — Service name.
- Description — Service description.
- Additional Required Components — Any other components required to run this service. You cannot enable this service unless this field is empty.
- Enable Service — Only enabled services are deployed. Disabling a service, effectively undeploys just that service while letting all other services in the application run as normal. This can be useful, for example when you wish to deploy an application that includes a service for which you don't have the required software.

Target Machines

- Remove from Selected Machines — Click to remove this service configuration from the selected machine(s).
- Add to Additional Machines — Adding services to additional machines is useful for fault tolerance. As a rule, it therefore does not make sense to run the same service on the same machine twice.
- Service Instance — Service instance from the selected machine. The service instance name includes the machine name.
- Software — The software required by this service instance.
- Deployment Status — Deployment status, as shown in the Configuration Builder

- **FT Weight** — The fault tolerance status and weight of the service instance. Is displayed only if **Run Fault Tolerant** is selected. See [Configuring Fault-Tolerant Engines on page 179](#) for an in-depth discussion of this topic.

FT Group Settings

Appears only if a TIBCO BusinessWorks process is assigned to additional machines. Note that TIBCO Adapter services cannot be assigned fault tolerant options.

- **Run Fault Tolerant** — If selected, the selected service instances will run in fault tolerant mode.
- **Heartbeat Interval (ms)** — The master engine of a fault-tolerant group broadcasts heartbeat messages to inform the other group members that it is still active. The heartbeat interval determines the time (in milliseconds) between heartbeat messages. In the event if one process engine fails, another engine detects the stop in the master's heartbeat and resumes operation in place of the other engine. All process starters are restarted on the secondary, and services are restarted to the state of their last checkpoint.
- **Activation Interval (ms)** — A standard TIBCO Rendezvous fault tolerant parameter, documented in chapter 15, *Developing Fault Tolerant Programs of the TIBCO Rendezvous Concepts*.

Secondary process engines track heartbeat messages sent from the master engine. This field specifies the amount of time to expire since the last heartbeat from the master before the secondary restarts the process starters and process engines.

The Heartbeat Interval should be smaller than the Preparation Interval, which should be smaller than the Activation interval. It is recommended that Activation Interval be slightly over 2 heartbeats.

- **Activation Delay (Preparation Interval)(ms)** — A standard TIBCO Rendezvous fault tolerant parameter, documented in the *TIBCO Rendezvous Concepts* chapter 15 *Developing Fault Tolerant Programs*).

When a master engine resumes operation, the secondary engine shuts down and returns to standby mode. For some situations, it may be necessary to ensure that the secondary engine has completely shut down before the master engine resumes operation.

This field is used to specify a delay before the master engine restarts. When the time since the last heartbeat from an active member exceeds this value, the ranking inactive member will receive a "hint" so that it can prepare for activation.

The Heartbeat Interval should be smaller than the Preparation Interval, which should be smaller than the Activation interval.

Monitoring Tab

Rulebases

Click **Add** to add an existing custom TIBCO Hawk rulebase. The rulebase must have been configured using TIBCO Hawk Display. See [Adding a Custom Rulebase to a Process or Service on page 170](#) for more information.

Events

Click **Add** to create an event. See [To Add an Event to a Service on page 174](#) for more information.

Failure Count

When an instance is down unexpectedly, the error count and last failure time are tracked. When the error count is greater or equal to the value set for Reset Failure Count, or if the value set for Reset Failure Interval expires (whichever comes first), the error count is reset to zero.

- **Reset Failure Count.** The value in this field defines how many restarts should be attempted before resetting the error counter to 0.

When an instance is down, the TIBCO Hawk agent will attempt to restart the instance the number of times specified in this field. If the instance restarts after the number of times specified, the event you have defined is triggered.



The event is *never* triggered if Reset Failure Count is set to 0.

- **Reset Failure Interval (seconds).** The value in this field defines how much time should expire before resetting the error counter to 0.

For example, if you define the following three events and set the Reset Failure Count to 5:

- Event 1, restart the instance and send an alert on the first failure.
- Event 2, restart the instance and send an email on the second failure
- Event 3, restart the instance and execute a command on subsequent failures.

On the first failure, the error count is 1, the instance is restarted and an alert is sent.

On the second failure, the error count is 2, the instance is restarted and email is sent.

On third failure, the error count is 3, the instance is restarted and the command you configured is executed.

On fourth failure, the error count is 4, instance is restarted and the command you configured is executed.

On fifth failure, the error count is 5 and then reset to 0. The instance is restarted and the command you configured is executed.

On sixth failure, the error count is 1, the instance is restarted and an alert is sent.

The cycle repeats.

If you do not want to receive alerts frequently, Reset Failure Count should be set with a high value. When error count is reset to 0, the last failure time is reset as well. The Reset Failure Interval takes effect only after the first failure occurs.

Advanced Tab

Click **Reset to Defaults** to use the defaults defined in the enterprise archive file.

Adapter SDK Properties

It allows you to change TIBCO Adapter SDK properties that were defined in the enterprise archive file.

Runtime Variables

It displays the service settable global variable for this service. You can change the global variable values as required. Be sure to use valid XML characters only.

TIBCO BusinessWorks Checkpoint Data Repository

If you wish to run TIBCO BusinessWorks using multiple engines in fault tolerant mode, you must specify a checkpoint data repository.

For true fault tolerance, you must store the data in a database.



You specify a JDBC Connection resource for the database to be used when you configure your project in TIBCO Designer. The database is then one of the available options on the pop-up menu.

TIBCO BusinessWorks Process Configurations

You can change the process configurations. For more information, see [Changing TIBCO BusinessWorks Process Configuration Properties on page 188](#).

The Flow Limit parameter appears in release 5.2 and later and only has meaning when deploying to a TIBCO BusinessWorks release 5.2 or later process engine. If you use Administrator 5.2 or later to deploy to a TIBCO BusinessWorks release 5.1.3 process engine, the parameter will display, but have no effect on the process engine.

Edit Service Instance Dialog

Fields can be edited if this dialog is invoked from the Configuration Builder pane. If invoked from the Deployed Configuration pane, the fields are read only.

The following tabs are available:

- [General Tab](#)
- [Server Settings Tab](#)
- [Graceful Shutdown Tab](#)

General Tab

The General tab displays the following information:

- Software that will run the used by the service instance.
- Machine on which this instance has been set up to run.
- Operating system used by this machine.
- Name of the service instance.



The name of the Service instance should not exceed 64 characters.

- Description for this service instance.
- Contact for this service instance.

Server Settings Tab

General

- **Start on Boot** — Specifies that the service instance is started when the machine on which the service instance is installed restarts. The service instance is not restarted when the administration server is restarted. Make sure that Hawk Agent is working.

For file based domain, service instance deployed with Start on Boot option starts only if the Administrator server is running.

For data-based domain, the Administrator server does not need to be running; the database must be up.

- **Enable Verbose Tracing** — Enables verbose tracing, in particular, for TIBCO BusinessWorks service instances.
- **Max Log File Size (KB)** — Specifies the maximum size (in Kilobytes) a log file can reach before the engine switches to the next log file.
- **Max Log File Count** — Specifies the maximum number of log files to use. When log files reach the size specified in the Max Log File Size field, the engine switches to the next log file. When the maximum number of log files have been written, the engine begins writing to the first log file again.
- **Thread Count**

Thread Count specifies the number of threads to use to execute process instances. The number of threads determines how many process instances can execute concurrently. Set the number of threads to a value that is appropriate for your operating system and physical machine configuration.

You should measure the available CPU and memory resources on your system under a typical processing load to determine if the default value of 8 threads is appropriate for your environment. For example, if engine throughput has reached a plateau, yet measurements show that CPU and memory are not fully utilized, increasing this value can have a positive effect on throughput. Typical numbers of worker threads range between 4 and 32. Specifying too low a value can cause higher memory use and lower engine throughput even though spare CPU resources exist. Specifying too high a value can cause CPU thrashing behavior, or an increase in latency caused by a large number of messages in the message queue.

Java

This pane is only available for Java applications.

- **Prepend to Classpath** — The items you supply here are prepended to your CLASSPATH environment variable. You can specify a Java code editor, or the jar file from a JNDI provider if you wish to use TIBCO BusinessWorks to receive and process JMS messages.
- **Append to Classpath** — The items you supply here are appended to your CLASSPATH environment variable. You can specify a Java code editor, or the jar file from a JNDI provider if you wish to use TIBCO BusinessWorks to receive and process JMS messages.
- **Initial Heap Size (MB)** — Initial size for the JVM used for the process engine. Default is 32 MB.
- **Maximum Heap Size (MB)** — Maximum size for the JVM used for the process engine. Default is 128 MB.
- **Java Thread Stack Size (KB)** — Size for the thread stack. Default is 128 KB.

NT Service

- **Run as NT Service** — Select to run this service as a Microsoft Windows Service. You can then manage the engine as you would any other service, and you can specify that it starts automatically when the machine reboots.



When running a service instance as an NT Service, you must use the TIBCO Administrator GUI to start and stop the instance. Manually starting or stopping the instance directly from the NT Service console is not supported.

- **Startup Type** — Choose one of the service startup types, Automatic, Manual, or Disabled.
- **Login As** — Specify the login account for the service, if any. The domain name must be specified. If the user is defined on the local machine, the domain is ".". For example, user jeff on the local machine would be specified as .\jeff.
- **Password** — Click set to define the password for that service, if any.

Graceful Shutdown Tab

This tab appears only if you have displayed this dialog box from a process. You can specify how a graceful shutdown occurs.

Kill Jobs Timeout

Specifies the maximum timeout in seconds after you click the Stop button that the process engine will wait for jobs to finish before shutting down the engine. A zero (0) value means 0 seconds, which effectively turns the graceful shutdown into an immediate shutdown.

Wait for Checkpoint

When selected, causes the process engine to wait for all jobs to finish after you click the Stop button (up to the maximum timeout) before shutting down the engine, rather than removing jobs at their next checkpoint.

Managing and Monitoring Process Engines and Service Instances

This chapter explains how to start and stop process engines and service instances. It also explains how to view the log files generated by the processes.

Topics

- [Process Engines and Service Instances Overview, page 213](#)
- [Starting or Stopping a Service Instance or Process Engine, page 215](#)
- [Viewing Log File Information, page 217](#)
- [Editing Process Engine Properties, page 219](#)
- [Viewing the TIBCO Administrator Audit Log, page 220](#)
- [All Service Instances Dialog, page 221](#)
- [View Service Instance Dialog, page 223](#)
- [View Service Instance: TIBCO Administrator Dialog, page 227](#)

Process Engines and Service Instances Overview

You can access process engines and service instances as follows:

To view all process engines and service instances for the administration domain, select the Application Management > All Service Instances console. It allows you to view and change the status of all instances running in the administration domain. The console displays the software for which there are running instances (for example, TIBCO Administrator, TIBCO Enterprise Message Service, or TIBCO BusinessWorks), and then allows you to view all instances for that software.

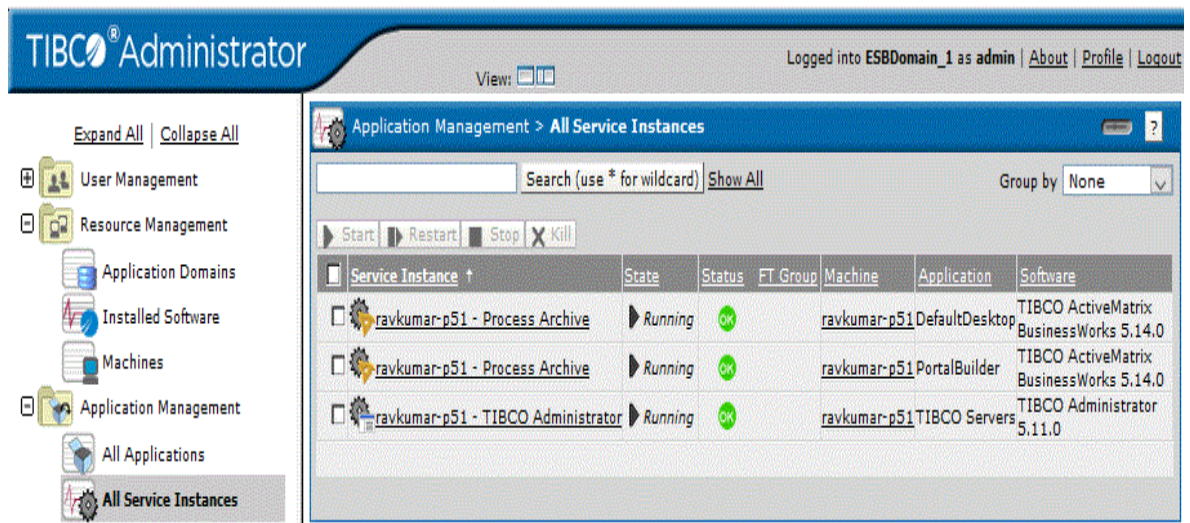
To view all process engines and service instances for an application, select the corresponding application, and then choose Service Instances. Only process engines and service instances for the that application are displayed.



When running a service instance on Microsoft Windows as an NT Service, you must use the TIBCO Administrator GUI to start and stop the instance. Manually starting or stopping the instance directly from the NT Service console is not supported.

The next diagram shows all service instances in an administration domain.

Figure 50 All Service Instances



The All Service Instances console is organized to show the software at top level. You can then select the software to display all associated service instances.

For example, if you created multiple deployments of a TIBCO BusinessWorks process engine, you will see one engine for each deployed process archive. You can select a TIBCO BusinessWorks process engine to view information about its corresponding jobs.

Starting or Stopping a Service Instance or Process Engine

After you have deployed an application, you can change the service instance or process engine state from different parts of the TIBCO Administrator GUI:

- In the Application Management > All Service Instances console, click the check box next to a service or process, and then choose the appropriate button, which becomes available.
- In the Application Management module, select the All Service Instances console and click the service instance name. In the window that is displayed, select the General tab. State has a clickable word next to it that allows you to start or stop service instances.
- In the Application Management module, select the Service Instances console for the application, click the check box next to a service instance, and then choose the appropriate button, which becomes available.
- You can start a service instance when it is deployed, or you can explicitly start services after deployment. See [Deploying an Application on page 141](#) for information about starting services when deploying.
- TIBCO Administrator itself is also listed as component software.



Shutting down the administration server is not recommended. You cannot start a stopped server from TIBCO Administrator.

You can, however, *restart* the administration server if you are using a database domain or have a secondary server defined in the domain: Select the server and choose **Restart**, which stops and restarts the server. In that case, a "page cannot be displayed" error results in the TIBCO Administrator GUI. You must invoke TIBCO Administrator GUI again and log in once more.

If you shut down the administration server, all currently running services and processes in the domain will continue to run. However, you can no longer monitor any project and the applications using Rendezvous or http as transport. You cannot restart any component in case of failure. In addition, some components load configuration information on demand which will fail if the administration server is not available.

To Start or Stop a Service Instance or Process Engine

1. Under Application Management, select **All Service Instances** or go directly to a specific application and select **All Service Instances**.
2. Select the service instances, process engines or both to start and click **Start**.

Click **Stop** or **Kill** to shut down the service instance or process engine. See [All Service Instances Dialog on page 221](#) for more information.

Viewing Log File Information

Tracing options are set in TIBCO Designer when configuring a service or process. You can view the tracing options for a process or service instance and set search parameters to display only a subset of a log file. You can also export the log to a file.

When you display the Trace tab, you specify the log file in this field to get information from and customize the number of lines to return in the Lines to fetch field. You can supply one or more search conditions to filter the amount of information to return.

- **Date/Time before/after** — Specify a date to filter. Use two dates to create a range.
- **Role** — Allows you to choose only traces to certain roles. Choose Info, Warning, Debug, Error, or specify a Custom role.

The role you choose depends on the role defined for the application. By default, Info, Warning, Error, and Debug are available. Custom roles may also be available if supported by the application.

- **Category** — Specifies a category. Items for that category are then sent to the trace. For example, if you choose Database, any database access or database errors are included.

Categories include, for example, Configuration, Application, Adapter, Database, TibRvComm, and XML. Custom categories may also be available if supported by the application.

- **Detail description** — Allows you to specify a detail description for which you want to display (or not) all log entries.

Make sure the process or service State indicates the process or service is running.

To View Tracing Results for a Process Engine or Service Instance

1. Under Application Management, select **All Service Instances** or go directly to a specific application and select **All Service Instances**.
2. Click the process engine or service instance name.
3. Click the Tracing tab.
4. Click **details** to drill-down view the entry.
5. To export an entry to a file, select the item(s) you wish to export.
6. Click **Export**.

7. Click **Done**.

Editing Process Engine Properties

You can edit active processes, process starters, process definitions and lock properties defined for a process engine.

To Edit Process Engine Activities Properties

1. Under Application Management, select **All Service Instances** or go directly to a specific application and select **All Service Instances**.
2. Click a process engine name.
3. Click the **BW Processes** tab.
4. Select an activity from the drop-down menu. The panel changes, depending on your selection.
5. Click **Done**.

See Also

See [BW Processes on page 223](#) for more information.

Viewing the TIBCO Administrator Audit Log

For TIBCO Administrator, you cannot configure tracing. You can, however, view the audit log, and filter it to better view the information you need.

In many cases, your browser does not display the complete log. In those cases, define a search condition as discussed in [Viewing Log File Information on page 217](#), and then click **Search**.

To View the Audit Log

1. Choose **Application Management > All Service Instances**.
2. Select **TIBCO Administrator**
3. Click the **Audit Log** tab.
4. Click the **Search** button to display the Administrator log.
5. Optionally, add a search condition and click **Search**.
6. Click **details** to drill-down view the entry.

To export an entry to a file, select the item(s) you wish to export.

1. Click **Export**.
2. Click **Done**.

All Service Instances Dialog

Search

Allows you to display only the items that match a search criteria.

Start

Starts the selected service instance or process engine.

Restart

Stops the selected service instance or process engine, and then starts it. This command should not be used for TIBCO Administrator unless it is using a database for domain storage or has a secondary server defined. See [Starting or Stopping a Service Instance or Process Engine on page 215](#) for details.

Stop

Stops the selected service instance or process engine. If graceful shutdown options are set for a process engine, the options are applied. Click a process engine name to access graceful shutdown options.

Kill

Forces an immediate shutdown of each selected service instance or process engine. If checkpoints or other graceful shutdown options are defined for a process engine, the options are ignored. Current jobs are terminated before given a chance to complete.

Group By


Determines how items in the display are grouped.

Instances List

- **Software** — Name of the installed TIBCO software that runs the application. The highest alert for that software is displayed in the left-most column.
- **Service Instance** — Displays the TIBCO BusinessWorks engine, adapter instance, JMS Server service, and so on. Click the component name for additional information.

- **State** — Stopped, Starting Up, Running, or Shutting Down. If the component belongs to an FT group, Standby is also an option.

Shutting down TIBCO Administrator is not recommended. Restarting is, however, an option.

- **Status** — Indicates the status for the application. The  icon indicates that the instance has lost contact with the endpoint Hawk Agent. Status cannot be determined.
- **FT Group** — Fault Tolerance group to which this component belongs, if any.
- **Machine** — The computer on which this component is running.

View Service Instance Dialog

The following tabs are available:

- [BW Processes](#)
- [General Tab](#)
- [Graceful Shutdown Tab](#)
- [Tracing Tab](#)
- [Security Tab](#)

BW Processes

This tab displays only for process engines.

Select Active Processes, Process Starter, Process Definitions or Locks.

Active Processes

Displays active process engines. As a rule, this includes process engines that are suspended or waiting. Examples include process engines that contain a Wait activity and are waiting in a loop. All other process engines usually complete before TIBCO Administrator is updated by auto-refresh and are therefore not displayed.

- **Add Search Condition** — You can add one or more search condition to narrow the display.
- **Export** — Click to export information about the selected process engine to a comma-separated file.
- **Kill** — Stops the selected process engine.
- **Suspend** — Suspends the selected process engine.
- **Resume** — Resumes the selected suspended process engine.

Process Starters

Displays all process starters in the process engine. You can then select individual process starters and enable or disable them. This can be useful, for example, if you wish to understand the performance impact of one of the process engines.

Process Definitions

Use the search field to limit the display. The * character can be used as a wildcard. Click a process engine to display the process definition details. The following information is displayed.

- Name — Activity for which throughput is displayed.
- Called Process — This field only shows information if you're running a process engine called by another process engine.
- Execution Count — Number of jobs in which this activity is currently participating.
- Elapsed Time (ms) — Time taken for all executions of the activity to complete.
- CPU time (ms) — CPU time used by this activity.
- Errors — Number of errors encountered for this activity.
- Status — Activity status.
- Function — Name of the activity resource.
- details — Click details for more detailed information about this particular activity.

Locks

Lock object shared configuration resources are used by Critical Section groups to ensure that only one process engine executes the activities within a Critical Section group at a time. The lock name, wait position, process id and requestor display. You can export lock information to a comma separated file, or kill a lock, if necessary. See the *TIBCO BusinessWorks Process Design Guide* for more information.

General Tab

General

Displays the following information about a process engine or service instance:

- Uptime for this component.
- Process ID for this component.
- Name of the process.
- Status of the component. If stopped, click **start** to start it. If running, click **stop** to stop it.

- Name of the machine on which this process engine or service instance is running.
- Name of the fault tolerant group, if any, to which this component belongs.

Statistics

This pane only displays for process engines.

- Created Processes — The processes created by the process engine.
- Suspended Processes — The processes currently suspended.
- Swapped Processes — The total number of times processes were swapped up to current.
- Queued Processes — The processes currently queued.
- Aborted Processes — The processes that were aborted.
- Completed Processes — The processes that were completed.
- Checkpointed Processes — The processes currently checkpointed.
- Total Execution (ms) — Total execution time for all processes. This refers to the total time the process was executing but does not include any wait times.
- Average Execution (ms) — Average duration for execution of a process.

Active Alerts

Displays information about the active alerts for this component.

- Date/Time — The date and time at which the alert occurred.
- Alert Level — The alert level set when the alert was created.
- Text — Description defined when creating the alert.

Graceful Shutdown Tab

Edit

Click to change the parameters under this tab.

Kill Jobs Timeout

Specifies the maximum timeout in seconds after you click the Stop button that the process engine will wait for jobs to finish before shutting down the engine. A zero (0) value means 0 seconds, which effectively turns the graceful shutdown into an immediate shutdown.

Wait for Checkpoint

When selected, causes the process engine to wait for all jobs to finish after you click the Stop button (up to the maximum timeout) before shutting down the engine, rather than removing jobs at their next checkpoint.

Tracing Tab

Allows you to view the trace logs for this application. You can create one or more search conditions to narrow the search scope.

To see the default log, leave **Where File is** *project.component.log* and click **Search**.

The log may grow quite large, and you are therefore encouraged to add one or more search conditions before you click **Search**. The number of lines displayed is governed by `tibcoadmin.monitor.traceLogMaxLines` in `tibcoadmin_domain.tra` and defaults to 1000.

Security Tab

See [Security Dialog on page 103](#) for details about these fields.

View Service Instance: TIBCO Administrator Dialog

The following tabs are available:

- [General Tab](#)
- [Audit Log Tab](#)
- [Security Tab](#)
- [Plug-Ins Tab](#)

General Tab

General

Displays detailed information about this component. See [All Service Instances Dialog on page 221](#) for information about the Stop and Restart commands.

Statistics

Displays detailed information about this component.

Active Alerts

Displays the active alerts for this component.

Audit Log Tab

Search

Allows you to search the default audit log, or a custom log. Click Add Search Condition to restrict the amount of information to return. To add multiple search restrictions, click the Add Search Condition button multiple times.

Security Tab

See [Security Dialog on page 103](#) for information about these fields.

Plug-Ins Tab

Plugin List

Each plugin that has been loaded into the administration domain is listed along with its name, specification titles and version, and implementation version.

Add

Displays a dialog where you can upload a plugin's WAR file.

Remove

Allows you to remove one or more selected plugins so they are no longer accessible in the administration domain.

Deploying, Starting and Monitoring an Adapter

This chapter provides a short tutorial that shows how to use TIBCO Administrator to deploy and start an application that contains an adapter configuration. Steps to deploy and start a TIBCO BusinessWorks process are similar and are discussed in Chapter 11, Deploying, Starting and Monitoring a TIBCO BusinessWorks Service.



Important Note: This chapter is meant to provide an introduction to the functionality, not comprehensive step-by-step instructions.

Topics

- [Prerequisites, page 230](#)
- [Opening the Project in TIBCO Designer, page 231](#)
- [Modifying the Adapter Service and Building the Archive, page 232](#)
- [Creating the Application in TIBCO Administrator, page 236](#)
- [Configuring the Application, page 238](#)
- [Deploying the Application, page 241](#)
- [Starting the Application, page 243](#)
- [Monitoring the Application, page 244](#)
- [Stopping the Service Instance, page 246](#)

Prerequisites

In this tutorial, you are stepped through preparing a project archive and deploying, starting, and monitoring the corresponding application. Complete the steps in this tutorial if you have TIBCO Adapter for Files installed and want to learn how to create an enterprise archive file in TIBCO Designer, and then deploy the file using TIBCO Administrator Enterprise Edition. If you don't have the adapter installed, read through the tutorial. The steps to create enterprise archive files and deploy them are the same for all adapters.

This tutorial uses the TIBCO Adapter for Files product and works with the delimitedReader example. The adapter reads a file and publishes the contents on a TIBCO Rendezvous subject.

Required Software

The following software must be installed to deploy the application. The software can be installed on one machine, or multiple machines. If installed on multiple machines, each machine must join the same administration domain.

- TIBCO Adapter for Files 5.1 or greater
- TIBCO Runtime Agent 5.3 or greater
- TIBCO Administrator Enterprise Edition 5.3 or greater

See Also

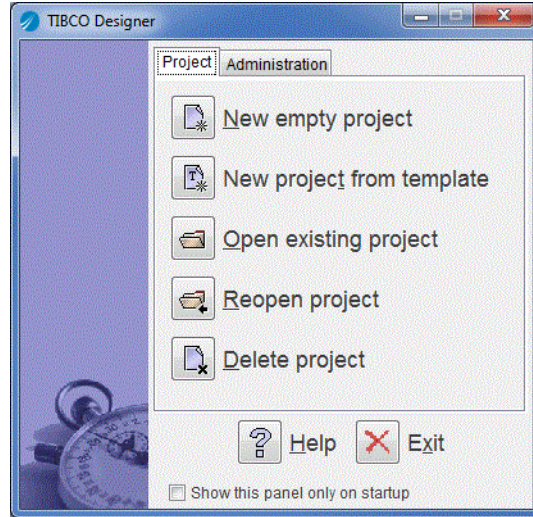
- See the *TIBCO Runtime Agent Domain Utility User's Guide* for information about using the TIBCO Domain Utility to add a machine to a domain.
- See the *TIBCO Designer User's Guide* for information about starting TIBCO Designer.
- See [Chapter 2, Starting TIBCO Administrator, on page 29](#).

Opening the Project in TIBCO Designer

TIBCO Adapter for Files includes an example file that is used in this example. The file is included in zip format and can be opened directly in TIBCO Designer.

1. Start TIBCO Designer and open a new empty project.

Figure 51 TIBCO Designer



2. Click **Cancel** in the Save Project dialog.
3. Click **Project > Import Full Project**.
4. Click the **Browse** button and import the project file, `TIBCO_HOME/adapter/adfiles/n.n/examples/Teak_TafRepoDefault_Files.dat`.
5. Click **OK**.
6. Select the **Replace existing global variables with those in import** and **Overwrite on name conflict** radio buttons in the Import Options dialog.
7. Click **Apply**.
8. Click **Project > Save** and save the project in a working directory.

Modifying the Adapter Service and Building the Archive

Whenever you wish to deploy a project, you must generate an enterprise archive file for it. In this section, you will build an enterprise archive file. You will also make two changes required for running this particular adapter service.

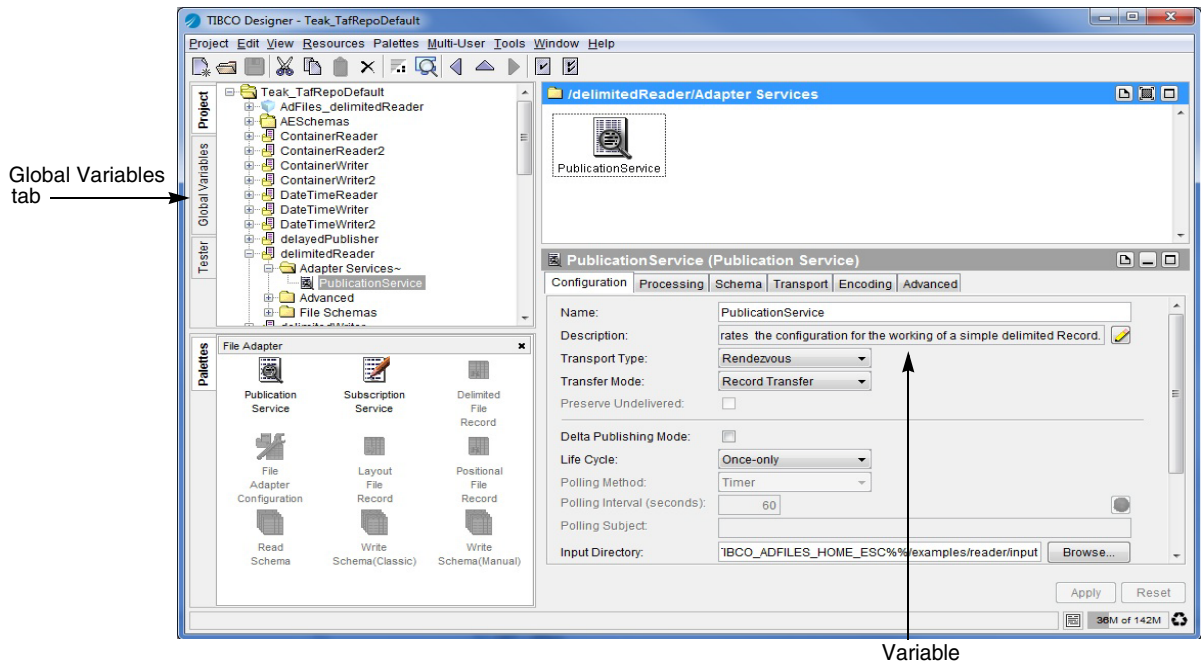
Modifying the Adapter Service

The `delimitedReader` example parses an input file and sends its contents in a series of TIBCO Rendezvous messages, and then stops. It uses a variable for the file name.

To successfully run and monitor the adapter from TIBCO Administrator, you must make the variable settable from TIBCO Administrator. It is also useful to make the service continuous. Follow the steps below:

1. With your project open, expand the **Teak_TafRepoDefault** folder.
2. Select **delimitedReader** in the project panel. Open it by clicking the plus (+) and open Adapter Services, and then select the Publication Service folder.
3. Select the **Processing** tab and notice that the value of the Working Directory is a variable.

Figure 52 Modifying the Adapter Service




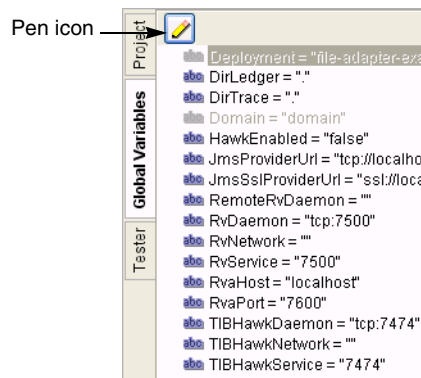

4. Click the **Global Variables** tab to the left of the project panel to display the Global Variables pane, and then select the  icon at the top.

Figure 53 Global Variables Tab



5. In the Global Variables editor, click the  icon to add a string variable. Triple-click in the Name field and type in the variable name (**TIBCO_ADFILES_HOME_ESC**). Select the **Deployment** and **Service** checkboxes, and then click **OK**.

Selecting those two checkboxes means that you can:

- Change the variable value for each deployment from TIBCO Administrator.
- Change the variable value for each service.

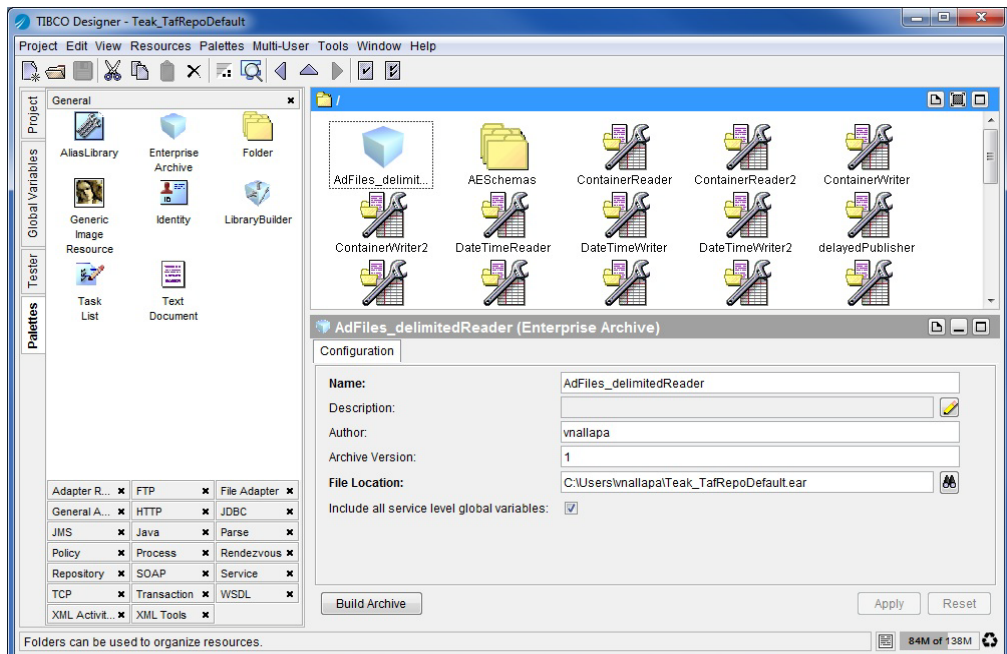
This would be useful, for example, if you had two publication services that used files from different directories.

Building the Archive

To build the enterprise archive file, follow these steps:

1. With the top-level (project) resource selected in the project tree, select the **General** palette in the palette panel and drag an Enterprise Archive resource into the design panel.
2. In the configuration panel, replace the default value in the Name field with **AdFiles_delimitedReader** and click **Apply**.

Figure 54 Building the Archive



3. In the project panel, select the **AdFiles_delimitedReader** enterprise archive. In the palette panel, select an Adapter Archive (from the Adapter Resources palette) and drag into the design panel.



4. In the Configuration panel, click the  icon next to the Adapter field and select the **delimitedReader** resource from the pop-up dialog, and then click **OK**.

Figure 55 Select a Resource



5. Click **Apply**.
6. Select the (top-level) **AdFiles_delimitedReader** enterprise archive in the project panel, click the **Build Archive** button in the bottom left corner of the configuration panel, and click **Yes** to save the project. A dialog appears indicating success.
7. Click **Project > Save** and then click **Project > Exit**.



Saving the project has no effect on the enterprise archive file you built. If you make additional changes to the project, you must rebuild the enterprise archive file.

Creating the Application in TIBCO Administrator

This section explains how to import an enterprise archive file and create a corresponding application using TIBCO Administrator.

Follow these steps:

1. Start TIBCO Administrator and log into the administration domain in which you wish to deploy the application.
2. Click the **Application Management** module, and then click the **New Folder** button. In the window that displays, in the Name field, type **File Adapter**. Click **Save**.



Creating folders for your application is recommended for complex applications. Using folders is not required.

3. In the left panel, click the **File Adapter** folder, and then click the **New Application** button.
4. Click the **Browse** button, and then select the enterprise archive file you created in [Opening the Project in TIBCO Designer on page 231](#).
5. Click **OK**.

If the file adapter is installed on multiple machines, you can choose an alternative target machine from the list before clicking Save.

Make sure an adapter is available as the Target in the bottom right corner. If it isn't, you must first install it on a machine in the administration domain.

Make sure the **Deploy on Save** option is *not* selected. In this example, you need to change a deployment option before you actually deploy, so you just want to create the application right now.

Figure 56 New Application Configuration

Application Management - File Adapter

New Application Configuration: AdFiles_delimitedReader

Save Cancel

General

Application Archive Change EAR File

Package Name	AdFiles_delimitedReader
Package Version	1
Package Description	
Package Creation Date	Mon Jan 28 00:07:00 PST 2008
Package Owner	Administrator


Application Parameters

Name	AdFiles_delimitedReader
Deployment Name	sample-AdFiles_delimit
Description	
Contact	
Max Deployment Revision	-1

Services

Quick Configure ☒

Deploy on Save ☐

Service	Description	Target
 delimitedReader.aar		clm - new adfiles 5.5.0.9

6. Click **Save**.

Configuring the Application

This section steps you through some application configuration tasks. These options are not required for all applications, but are required for this example. The following configuration tasks are discussed:

- Setting the Variable for the Service
- Setting Application Options

Setting the Variable for the Service

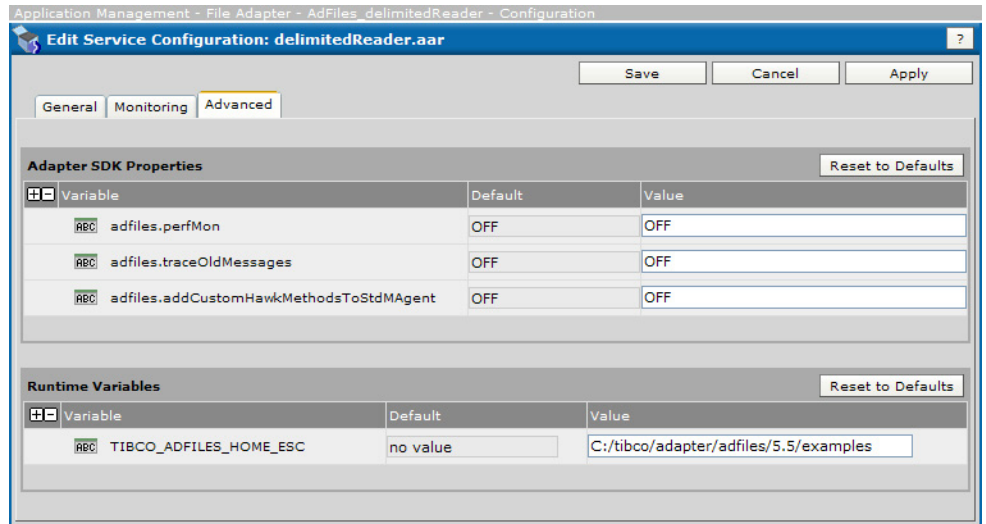
This section explains how to set the variable that specifies where the adapter will look for the file to parse. You're setting this variable for the service. Follow these steps:

1. In the left pane of TIBCO Administrator, choose **Application Management > File Adapter > AdFiles_delimitedReader > Configuration**.
2. In the Configuration Builder panel, click **delimitedReader.aar** and click the **Advanced** tab.
3. In the Runtime Variables pane, you can now specify a value for the global variable you predefined in TIBCO Designer. Supply the location of the examples directory for the file adapter, (*TIBCO_HOME/adapter/adfiles/version/examples*) then click **Save**.



You **must** use forward slashes for the directory name, even if the directory is on a Microsoft Windows machine.

Figure 57 Setting Variable for the Service



Setting Application Options

This section explains how to change certain options at the application level. Follow these steps:

1. In the left pane of TIBCO Administrator, choose **Application Management > File Adapter > AdFiles_delimitedReader > Configuration**.
2. In the Configuration Builder panel, click **AdFiles_delimitedReader**.
3. Click the **Advanced** tab and scroll down to the TIBCO BusinessWorks and Adapters Deployment Repository Instance panel.
4. In Transport, select **rv**.
5. Change the Discovery Timeout to **15**, and then click **Save**.

Figure 58 Setting Application Option

The screenshot shows a configuration window titled "TIBCO BusinessWorks and Adapters Deployment Repository Instance". It contains a list of configuration properties, each with a label and a corresponding input field. The properties and their values are as follows:

Property	Value
Transport	rv
Server Name	sample
Instance Name	sample-AdFiles_delimitedReader
User Name	cephas
Password	***** change...
Timeout	600
Service	7500
Network	
Daemon	tcp:7500
Discovery Timeout	15
Regional Subject	
Operation Retry	0

At the bottom left of the window, there is a button labeled "Preview URL".

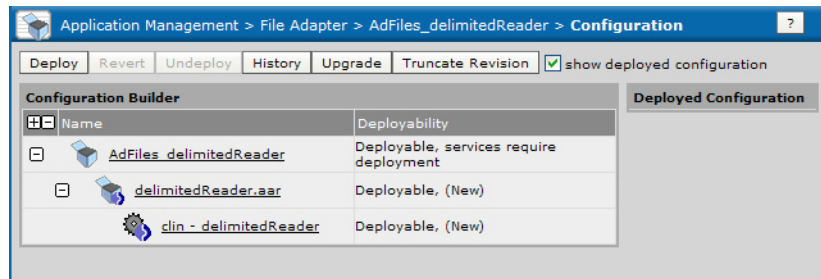
The Discovery Timeout property determines how long an application waits when initially trying to connect to the administration server. See [Changing Global Variables at Deployment on page 164](#) for a complete set of deployment configuration options.

Deploying the Application

When configuration is complete, you deploy the application. Follow these steps:

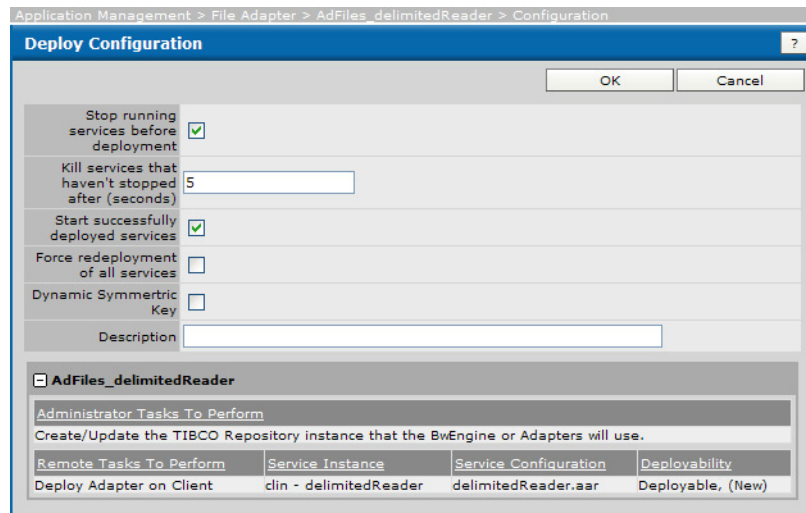
1. In the left pane, click **Application Management > File Adapter > AdFiles_delimitedReader > Configuration**.
2. In the Configuration Builder click the **Deploy** button.

Figure 59 Deploying the Application




3. Enter a description for this deployment. A description is useful should you later wish to revert to this configuration. Click **OK**.

Figure 60 Enter Description for Deployment



4. TIBCO Administrator returns to the Configuration console and updates the panel with information in the Deployed Configuration pane.

Figure 61 Configuration Console

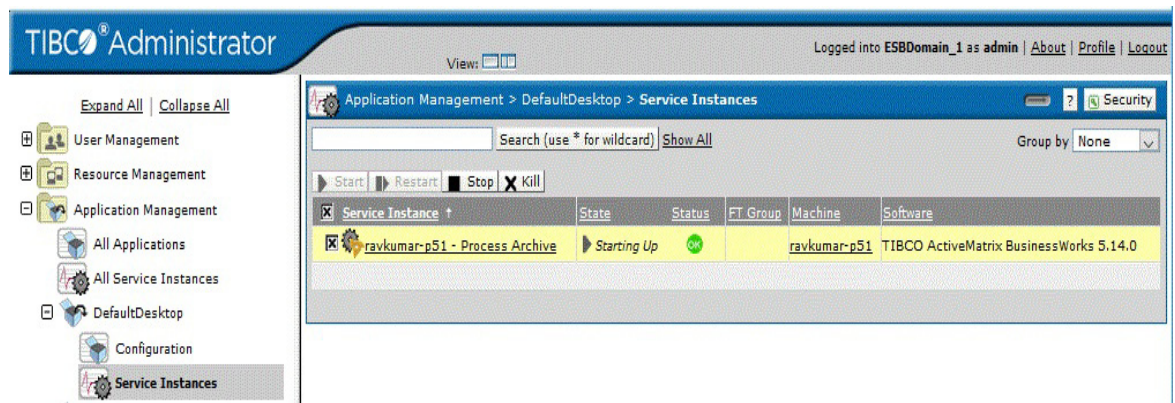


Starting the Application

This section explains how to start the adapter from TIBCO Administrator.

1. In the left pane, click **Application Management > File Adapter > AdFiles_delimitedReader > Service Instances**.
2. In the Service Instances panel, select the check box next to the `adfiles` service instance.
3. Click the **Start** button.

Figure 62 Starting the Application



Monitoring the Application

TIBCO Administrator creates a log for each running application. The actual file is stored in the application log directory. For example, `TIBCO_HOME/tra/domain/domain/application/logs`.

All log files use the following format: `application.application-component.log`, where `application` is the application name and `component` is the name of the application's adapter archive.



Do *not* change the names of the log files.

To view the log from TIBCO Administrator, follow these steps:

1. Choose **Application Management > File Adapter > AdFiles_delimitedReader > Service Instances** and select the delimitedReader Service Instance.
2. Select the Tracing tab and select the trace file for this service from the popup, and then click **Search**.

Figure 63 Monitoring the Application

Search

File is AdFiles_delimitedReader.AdFiles_delimitedReader-delimitedReader.log

Lines to fetch less than 1000

Add Search Condition

Search

Export

<input type="checkbox"/>	Date/Time	Role	Message Code	Description	
<input type="checkbox"/>	25-Mar-08 20:23:29 PDT America/Los_Angeles	Info	AEFA-000059	Scanning Input Directory C:/tibco/adapter/adfiles/5.5/examples/reader/input	details
<input type="checkbox"/>	25-Mar-08 20:23:29 PDT America/Los_Angeles	Info	AEFA-000058	TIBCO Adapter for Files successfully initialized	details
<input type="checkbox"/>	25-Mar-08 20:23:29 PDT America/Los_Angeles	Info	AEFA-000060	Processing input file delimited.txt in Input Directory C:/tibco/adapter/adfiles/5.5/examples/reader/input	details
<input type="checkbox"/>	25-Mar-08 20:23:29 PDT America/Los_Angeles	Info	AEFA-000065	Message containing class /tibco/public/class/ae/FileAdapter/wire/delimitedReader/Order published on subject delimitedReader (message is from file C:/tibco/adapter/adfiles/5.5/examples/reader/w...	details
<input type="checkbox"/>	25-Mar-08 20:23:29 PDT America/Los_Angeles	Info	AEFA-000065	Message containing class /tibco/public/class/ae/FileAdapter/wire/delimitedReader/Order published on subject delimitedReader (message is from file C:/tibco/adapter/adfiles/5.5/examples/reader/w...	details
<input type="checkbox"/>	25-Mar-08 20:23:29 PDT America/Los_Angeles	Info	AEFA-000064	File C:/tibco/adapter/adfiles/5.5/examples/reader/wip\delimited.txt has been processed, all lines were interpreted	details
<input type="checkbox"/>	25-Mar-08 20:23:29 PDT America/Los_Angeles	Info	AEFA-000079	File C:/tibco/adapter/adfiles/5.5/examples/reader/wip\delimited.txt has been parsed. Total: 8 lines, Error: 0 lines	details
<input type="checkbox"/>	25-Mar-08 20:23:29 PDT America/Los_Angeles	Info	AEFA-000063	Publication of file C:/tibco/adapter/adfiles/5.5/examples/reader/wip\delimited.txt is finished.	details

TIBCO Administrator displays a log for the activities of the service instance.



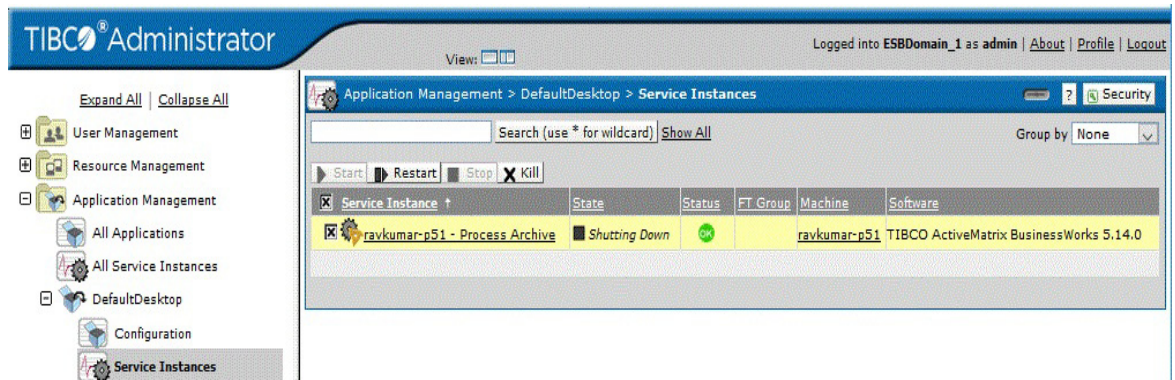
You can also configure monitoring events and TIBCO Hawk rulebases for monitoring. See [Adding a Custom Rulebase to a Process or Service on page 170](#).

Stopping the Service Instance

This section explains how to stop the running adapter service from TIBCO Administrator.

1. In the left pane, expand **Application Management > File Adapter AdFiles_delimitedReader > Service Instances**.
2. In the Service Instances panel, select the **adfiles** service instance.
3. Click the **Stop** button.

Figure 64 Stopping the Service Instance



Deploying, Starting and Monitoring a TIBCO BusinessWorks Service

This chapter provides a short tutorial that shows how to use TIBCO Administrator to deploy and start an application that contains a TIBCO BusinessWorks service.



Important Note: This chapter is meant to provide an introduction to the functionality, not comprehensive step-by-step instructions.

Topics

- [Overview, page 248](#)
- [Starting TIBCO Designer and Saving Your Project, page 250](#)
- [Creating the FileTest Process, page 251](#)
- [Testing the FileTest Process, page 257](#)
- [Creating the Enterprise Archive File, page 259](#)
- [Creating and Deploying the Application in TIBCO Administrator, page 260](#)
- [Starting the Application, page 261](#)
- [Monitoring the Application, page 262](#)
- [Stopping the Application, page 266](#)

Overview

Performing a deployment consists of a number of basic tasks:

TIBCO Designer Tasks

TIBCO Designer performs the following tasks (discussed in detail later):

1. [Starting TIBCO Designer and Saving Your Project on page 250](#)
2. [Creating the FileTest Process on page 251](#)
3. [Testing the FileTest Process on page 257](#)
4. [Creating the Enterprise Archive File on page 259](#)

See *TIBCO Designer User's Guide* for additional information about using TIBCO Designer to create a project and to build an enterprise archive file.

TIBCO Administrator Tasks

After creating the enterprise archive file in TIBCO Designer, you perform the following tasks (discussed in detail later) in TIBCO Administrator:

1. [Creating the Application in TIBCO Administrator on page 236](#)
2. [Starting the Application on page 261](#)
3. [Monitoring the Application on page 262](#)
4. [Stopping the Application on page 266](#)

Overview of Example Process



This sample project is created in the QuickStart tutorial that is included with the TIBCO BusinessWorks documentation set. If you went through that tutorial, you can use the project you created there.

The project monitors a directory for a specific file. When the file changes, a new file is created that contains the contents of the original file plus the time the change was made to the original file. The new file is named after the change that occurred (`create.txt`, `modify.txt`, or `remove.txt`). If you modify the file multiple times, the new file overwrites the existing `modify.txt`.

With this simple project, you will perform many of the same tasks that are required for larger, more complex projects. This tutorial is not intended to illustrate every aspect of TIBCO BusinessWorks, so only a small subset of the available activities will be used.

The tutorial steps you through the following tasks:

- [Starting TIBCO Designer and Saving Your Project](#)
- [Creating the FileTest Process](#)
- [Testing the FileTest Process](#)

Prerequisites

To perform the tasks in this tutorial, you must have installed and configured the TIBCO BusinessWorks and TIBCO Administrator software properly.

1. Install all components of TIBCO Runtime Agent *<version>* on your system.
2. Install all components of TIBCO BusinessWorks 5.12 or later on your system.
3. Install all components of TIBCO Administrator *<version>* on your system.

Starting TIBCO Designer and Saving Your Project

To start TIBCO Designer and save your server-based project, follow these steps:

1. From the Start menu, choose **All Programs > TIBCO > TIBCO Designer <version> > Designer <version>**.

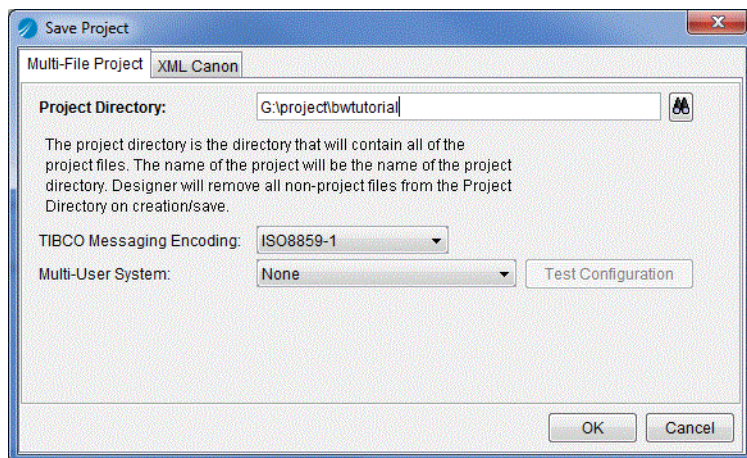
The TIBCO Designer startup window appears.

2. Choose **New empty project**.

The TIBCO Designer main window appears, with the Save Project dialog on top.

3. In the Save Project dialog, make sure the Multi-File Project tab is selected (the default).
4. In the Project Directory field, click the **Browse** button to locate the directory in which you wish to save the project. Click **OK**.

Figure 65 Save Project Dialog



5. Click **OK** to create the project.

You are now ready to create a process definition using TIBCO Designer.



For some introductory information on how to use TIBCO Designer, see the *TIBCO Designer User's Guide*.

Creating the FileTest Process

This section guides you through creating a simple process definition. The process, will be named FileTest. It:

1. Polls a directory for a specified file
2. Writes a new file to the same directory each time the file changes.
 - The new file's name includes the type of change that occurred to the original file (create, modify, or remove).
 - The new file's content is the same as the polled file's content, but the time of the change in the file is appended to the end of the file. The time of the file change is represented as the number of milliseconds since January 1, 1970.

The tutorial uses a variable file name to illustrate how to use the TIBCO BusinessWorks mapping capabilities.

Create the Example Directory and File


The File Poller activity requires an input file.

1. Create a directory for the input file. For example, `c:\projects\bwtest`.
2. Create a file in the directory named **PolledFile.txt**.
3. Add the following text to the file:

The cherry blossoms are beautiful.

Create the FileTest Process

1. In TIBCO Designer, select the project name (**bwtutorial**) in the project panel.
2. In the palette panel, select the **Process** palette.

If no palettes are in the palette panel, click the **Switch Palette Mode**  icon to display the palettes.
3. From the palette panel, drag a **Process Definition** into the design panel.
4. In the configuration panel, type the name **FileTest** in the **Name** field to rename the process. Then click **Apply**.
5. Save your project by choosing **Project > Save** from the menu.

Add Activities to the Process

1. Select the **FileTest** process in the project tree.

The Start and End activities should be displayed in the design panel.

2. Find the **File** palette in the palette panel and select it.



If the File palette is not one of the available palettes, click **Palettes > Browse** to display the Palette Browser. Type **File** in the Filter string/pattern field and click **Filter**. Select the File palette and click **Close**.

3. Drag a **File Poller** activity into the design panel (but not on top of the Start activity).

The Start activity is replaced by the File Poller activity.


4. With the File Poller still selected, enter the following values in the configuration panel:
 - a. Type **TestFilePoller** in the **Name** field.
 - b. Click the  button to the right of the File Name field and select the file to poll. Use the Select File dialog to locate the directory and file you created earlier (C:\projects\bwtest\PolledFile.txt).
 - c. Select **text** in the drop-down list in the Content as field, and select the appropriate encoding for your operating system in the Encoding field. Other fields should use defaults.

Figure 66 Adding Activities to the Process


The screenshot shows the 'TestFilePoller (File Poller)' configuration window. It has tabs for Configuration, Advanced, Misc, Output, and Error Output. The Configuration tab is selected. The fields are as follows:

Name:	TestFilePoller
Description:	
File Name:	C:\projects\bwtest\PolledFile.txt
Polling Interval (sec):	5
Include Existing Files:	<input type="checkbox"/>
Exclude File Content:	<input type="checkbox"/>
Content as:	text
Encoding:	ASCII

- d. Click **Apply**.
5. Drag a **Create File** activity into the design panel and place it to the right of TestFilePoller.

6. Name the Create File activity **CreateTestFile**, and click the **Overwrite** checkbox, and then click the **Apply** button.

Figure 67 Enter the Details in Configuration Tab



7. Drag a **Write File** activity into the design panel and place it to the right of the CreateTestFile activity (before the End activity).
8. Name the Write File activity **WriteToTestFile**, and then click the **Apply** button.

Figure 68 Activity Icons



9. Click **View** in the menu bar and select **Connect**.
10. Drag a transition from TestFilePoller to CreateTestFile.
11. Drag a transition from CreateTestFile to WriteToTestFile.
12. Drag a transition from WriteToTestFile to End.

The result should appear as follows:



13. Choose **Project > Save** from the menu.


The process now includes appropriately connected activities. However, no information about the name and content of the file to be created is included. To set those, you use the TIBCO Designer mapping facilities.

The goal in this exercise is to create a file that has the name `Filechange_type.txt`, where `change_type` is the type of change that occurred to the original file (create, modify, or remove). The content of the changed file is then written to the new file, and the time the change in the file occurred is appended to the end of the contents. For example, if the text in the changed file is **The cherry blossoms are beautiful.** the content of the new file will be **The cherry blossoms are beautiful. 1017163931036.**

To map the data flow between activities, follow these steps:

1. Click **View** in the menu bar and select **Select**.
2. Select the **CreateTestFile** activity and choose the **Input** tab in the configuration panel. Expand the **CreateActivityInputClass** item in the activity input tree by clicking on the plus sign (+) to the left of the item.


Figure 69 Mapping Data Flow between Activities



Each activity's output is available to subsequent activities in the process definition. You can use data from previous activities to specify the input of the current activity. For example, you can use the content of the original text file as the content of the new text file.

The left panel of the Input tab contains a list of data from all activities preceding the current activity in the process diagram. Each activity's name appears with a dollar sign (\$) in front of it to indicate that this is a process variable.


The right panel of the Input tab lists the current activity's input. Input fields in red indicate an error in the data specified for the field. In this case, the `fileName` field is required. Because no value has been supplied yet, the field is displayed in red.

3. Click the plus (+) sign next to the process variable `$TestFilePoller` in the Process Data panel and expand the schema tree containing the output of the `TestFilePoller` process starter.
4. In the right panel, expand the tree and select the `fileName` field in the Activity Input pane, and then click the  icon.

5. In the XPath Formula Builder dialog that appears, follow these steps:
 - a. Select the **Functions** tab, expand the **String** folder, select **concat**, and drag it into the XPath Formula panel.
The display in the right panel changes to display a concat XPath expression.
 - b. Replace `<< string1 >>` with `"C:\projects\bwtest\File"` (include the quotes).
 - c. In the left panel, select the Data tab, choose the `$TestFilePoller/EventSourceOutputTextClass/action` item and drag it over `<< string2 >>`. A red box appears over `<< string2 >>` indicating you can release the data over this item and the correct XPath expression will appear.
 - d. Add a comma after `$TestFilePoller/EventSourceOutputTextClass/action`, and then add the string `".txt"` (include the quotes).
The expression should look like the following:

```
concat("c:\tibco\test\File",
$TestFilePoller/EventSourceOutputTextClass/action, ".txt")
```
6. Click the **Apply** button to accept the formula, and then click **Close**. Click the **Apply** button in the configuration panel of the activity.
7. Select the **WriteToTestFile** activity, and then click the **Input** tab in the configuration panel.
8. Map the data as follows:
 - a. Select `$CreateTestFile/CreateActivityOutputClass/fileInfo/fullName` and drag it to the `fileName` item in the Activity Input pane.
You do not need to use the XPath Formula Builder to map this item because you do not need to use XPath functions. The content of the field is exactly the same as the value of the

\$CreateTestFile/CreateActivityOutputClass/fileInfo/fullName process variable.

- b. In the right panel, select the **textContent** field and click the  icon.
- c. In XPath Formula Builder, drag a concat function into the XPath formula panel.
- d. Add " ", between `<<string1>>`, and `<<string2>>` so that there is a space between the two strings in the concat function.
- e. Click the Data tab and drag **\$TestFilePoller/EventSourceOutputTextClass/fileContent/textContent** over `<< string1 >>`.
- f. Drag **\$TestFilePoller/EventSourceOutputTextClass/timeOccurred** over `<< string2 >>`.

The formula should look like this:


```
concat($TestFilePoller/EventSourceOutputTextClass/
fileContent/textContent," ",$TestFilePoller/
EventSourceOutputTextClass/timeOccurred )
```

- g. Click the **Apply** button to accept the formula and click **Close**.
9. Click the **Apply** button on the activity's configuration panel, and then choose **Project > Save** to save your project.


You are now ready to test the project.

Testing the FileTest Process

You can test the FileTest process directly from TIBCO Designer. This allows you to make sure the process works correctly before you deploy it. This step is not strictly required before you deploy but highly recommended. Follow these steps:

1. Type **Alt+F8** to open the Set BreakPoints dialog.
2. In the window that appears, select **Select All**, and then click **OK**.
Breakpoints allow you to step through the process. Stepping helps you see what happens when each activity executes.
3. Click the **Tester** tab to the left of the project panel. The test panel replaces the project tree.
4. Click the  button.
5. In the process selection window that appears, the FileTest process is selected by default. Click **Load and Start Current**.

The process is now in Test mode.

6. Make a change to the text in `C:\projects\bwtest\PolledFile.txt` to start the process.
7. Once the TestFilePoller process starter is highlighted (indicating a process has started), click the  icon to step through the process.



The icon has two boxes, not three. You can see the name of the icon if you allow the cursor to rest on it without moving the mouse.


TIBCO ActiveMatrix BusinessWorks creates an output file in the same directory named `Filemodify.txt` after you have stepped into the WriteToTest activity and writes the appropriate text to the file.

The text should be the text of `PolledFile.txt` and, after a space, the time, in milliseconds, since January 1, 1970.

You can click on each activity in the process definition as you step through it. If you click on the Process Data or Output tabs for the activity, you will see the actual process data and output of the activity as the process executes.

8. Next, delete **PolledFile.txt**.
9. Step through the process once more.

TIBCO Administrator creates a file `Fileremove.txt`. The content of `Fileremove.txt` is just the time of modification, because the polled file no longer exists.

10. Click the  icon to return to design mode. In design mode, you can now prepare your enterprise archive.

For more information on using test mode, see the *TIBCO BusinessWorks Process Design Guide*.


Creating the Enterprise Archive File

Before you can deploy a project, you must create an enterprise archive file in TIBCO Designer. Follow these steps:

1. If your project is not currently open, choose **Open Existing Project** in the startup screen and select the project you wish to deploy.



If you worked with the project recently, you can also choose **Reopen Project**.

2. Select the top-level (**bwtutorial**) folder and, from the palette panel, drag an Enterprise Archive resource (from the General palette) into the design panel.
Notice that the name is the same as the project name.
3. In File Location, use the default, or click browse and provide a location and filename.
4. With the tutorial archive selected in the project panel, drag a Process Archive (found in the Process palette) into the design panel.
5. In the configuration panel
 - a. Change the name to **FileActivityTest** and click **Apply**.
 - b. Select the **Processes** tab and click the  icon.
 - c. Select the FileTest process you created earlier, click **OK**, and then click **Apply**.
6. Select the **bwtutorial** archive in the project panel and click the **Build Archive** button in the bottom left corner of the configuration panel.
7. Click **Yes** to save the project.

Creating and Deploying the Application in TIBCO Administrator

This section explains how to use TIBCO Administrator to import an enterprise archive file and create a corresponding application.

Follow these steps:

1. Start TIBCO Administrator and log into the administration domain in which you wish to deploy the application.
2. Click the **Application Management** module, and then click the **New Folder** button.
3. In Name, type **Timer Application**.
4. Click **Save**.
5. Double-click the Timer Application folder, and then click the **New Application** button.
6. Click the **Browse** button to select the enterprise archive file you created in [Creating the Enterprise Archive File on page 259](#).
7. In the dialog that appears, select the **Deploy on Save** checkbox, and then click the **Save** button.



If you choose Deploy on Save, TIBCO Administrator uses the parameters specified in the project file and the default machine that registered the software in TIBCO Administrator.

This example does not require further customization. For other cases, you may decide not to choose Deploy on Save so you can first configure the application.

The next dialog displays the application. In the left panel, expand the application and click Configuration. This displays the Configuration Builder and Deployed Configuration panels with the consoles created for the deployment. The Configuration Builder panel on the left allows you to customize the application configuration. The Deployed Configuration panel shows the deployed applications.

Starting the Application

This section gives an overview of starting an application.

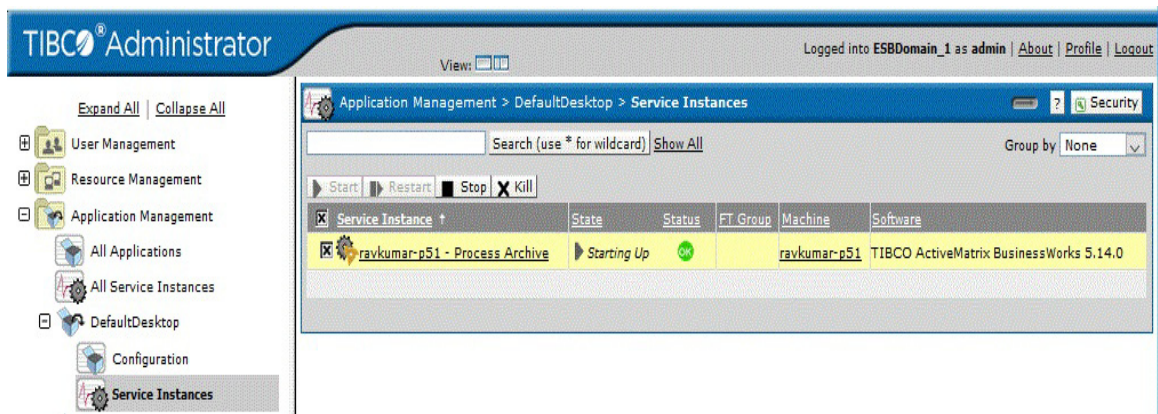


If you deployed using Deploy on Save, the tutorial application is actually started by default.

To start service instances, follow these steps:

1. In TIBCO Administrator, select the application in the left panel, and then click Service Instances.

Figure 70 Starting the Application



2. In the Service Instances console, click the check box next to the Service Instance (named after the machine and the process archive) and click **Start**. The State column changes to first show **Starting Up**, and then **Running**.

Monitoring the Application

Monitoring an application can be done in two ways, discussed in this section:

- Viewing Default Monitoring Information
- Specifying a Custom Alert

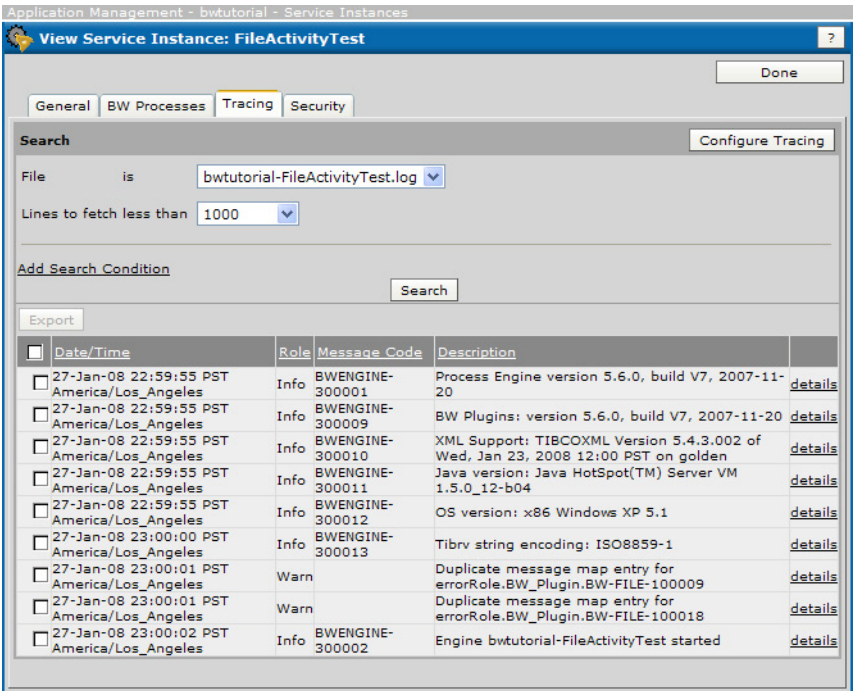
Viewing Default Monitoring Information

Some monitoring information for the application is available by default. To view the log from TIBCO Administrator, follow these steps:

1. Choose **Application Management > Timer Application > bwtutorial > Service Instances** and select the *machine-FileActivitiesTest* service instance, and then the Tracing tab.
2. In the Search box, make sure that the log for your application (in this case *bwtutorial-FileActivityTest.log*) is selected, and then click **Search**.

TIBCO Administrator displays information about the instance, which includes starting, termination, and any errors that occurred.

Figure 71 Viewing Default Monitoring Information




3. To make the default monitoring information more detailed, click the **Configure Tracing** button.



The service instance must be running or the button will be disabled.

4. In the window that appears, click **All Activities** to have execution of each activity included in the log; click **All Starters** to have execution of all starters included, and then click **Save**.

Figure 72 Configure Tracing



- When you return to the log, you will find that information about the individual activities and the starter are now included.

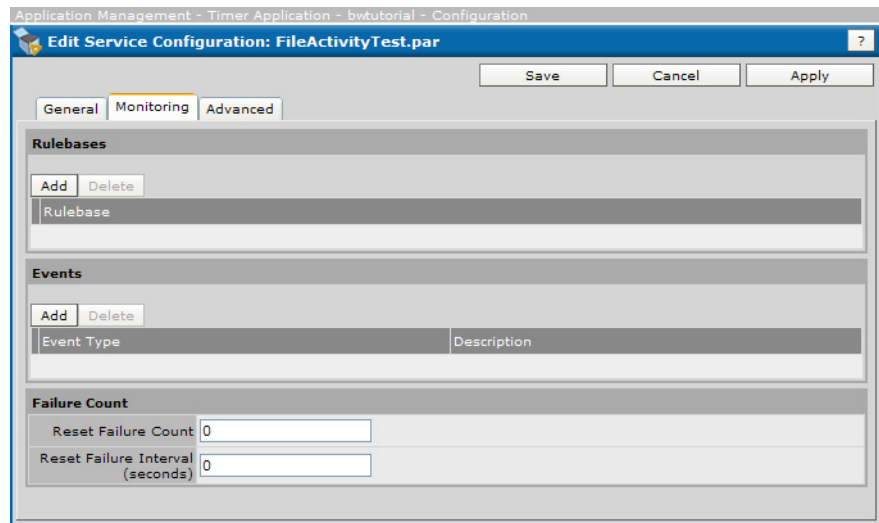
Specifying a Custom Alert

In addition to tracing, using TIBCO Administrator you can specify that you wish to be alerted if certain conditions are met. This section gives one example. A detailed discussion of tracing is included in [Adding a Custom Rulebase to a Process or Service on page 170](#).

To specify a custom alert, follow these steps:

- Choose **Application Management > Timer Application > bwtutorial > Configuration**.
- In the Configuration Builder panel, expand the bwtutorial application and click **FileActivityTest.par**, and then click the **Monitoring** tab.

Figure 73 Specifying a Custom Alert



- In the Events pane, click **Add**.

4. In the dialog that appears, make the following changes (shown in the next figure).
 - a. In the General pane, select **First Component Failure** in the Event Type drop-down list.
 - b. In the Alert pane, select the **Generate Alert** checkbox.
 - c. Select the **All Occurrences** radio button and select **High** in the Level drop-down list.
 - d. In the Message field, type **First component failure - FileActivityTest**.
 - e. Click **OK**, and then click **Save**.

TIBCO Administrator will now display a high-level alert with the message upon first failure of this process. If you wanted, you could also have an email sent in the event of component failure.

Figure 74 Add Event

Application Management - Timer Application - bwTutorial - Configuration
Edit Service Configuration: FileActivityTest.par

Add Event ?

OK Cancel

General

Event Type: First Component Failure

Restart Service Instance: ☐

Description:

Alert

Generate Alert: ☐

Generate Alerts For: ☒ First Occurrence ☐ All Occurrences

Level: High

Message: First component failure - FileActivityTest

5. When you are returned to the Configuration Builder notice that it indicates that services require deployment. Click **Deploy** and the changes will take effect.

Stopping the Application

To stop the application, follow these steps:

1. In the left panel of TIBCO Administrator, select the Application Management module.
2. Select either the All Service Instances panel, or choose **Application Management > Timer Application > bwtutorial > Service Instances**.
3. Select the check box next to the process engine you started earlier, and then click **Stop**.

TIBCO Product Documentation and Support Services

For information about this product, you can read the documentation, contact TIBCO Support, and join the TIBCO Community.

How to Access TIBCO Documentation

Documentation for TIBCO products is available on the [TIBCO Product Documentation](#) website, mainly in HTML and PDF formats.

The [TIBCO Product Documentation](#) website is updated frequently and is more current than any other documentation included with the product.

Product-Specific Documentation

The following documentation for this product is available on the [TIBCO Administrator Enterprise Edition](#) page.

- *TIBCO Administrator™ Release Notes*
- *TIBCO Administrator™ Installation*
- *TIBCO Administrator™ Server Configuration Guide*
- *TIBCO Administrator™ User's Guide*

How to Contact TIBCO Support

Get an overview of [TIBCO Support](#). You can contact TIBCO Support in the following ways:

- For accessing the Support Knowledge Base and getting personalized content about products you are interested in, visit the [TIBCO Support](#) website.
- For creating a Support case, you must have a valid maintenance or support contract with TIBCO. You also need a user name and password to log in to [TIBCO Support](#) website. If you do not have a user name, you can request one by clicking **Register** on the website.

How to Join TIBCO Community

TIBCO Community is the official channel for TIBCO customers, partners, and employee subject matter experts to share and access their collective experience. TIBCO Community offers access to Q&A forums, product wikis, and best practices. It also offers access to extensions, adapters, solution accelerators, and tools that extend and enable customers to gain full value from TIBCO products. In addition, users can submit and vote on feature requests from within the [TIBCO Ideas Portal](#). For a free registration, visit [TIBCO Community](#).

Legal and Third-Party Notices

SOME TIBCO SOFTWARE EMBEDS OR BUNDLES OTHER TIBCO SOFTWARE. USE OF SUCH EMBEDDED OR BUNDLED TIBCO SOFTWARE IS SOLELY TO ENABLE THE FUNCTIONALITY (OR PROVIDE LIMITED ADD-ON FUNCTIONALITY) OF THE LICENSED TIBCO SOFTWARE. THE EMBEDDED OR BUNDLED SOFTWARE IS NOT LICENSED TO BE USED OR ACCESSED BY ANY OTHER TIBCO SOFTWARE OR FOR ANY OTHER PURPOSE.

USE OF TIBCO SOFTWARE AND THIS DOCUMENT IS SUBJECT TO THE TERMS AND CONDITIONS OF A LICENSE AGREEMENT FOUND IN EITHER A SEPARATELY EXECUTED SOFTWARE LICENSE AGREEMENT, OR, IF THERE IS NO SUCH SEPARATE AGREEMENT, THE CLICKWRAP END USER LICENSE AGREEMENT WHICH IS DISPLAYED DURING DOWNLOAD OR INSTALLATION OF THE SOFTWARE (AND WHICH IS DUPLICATED IN THE LICENSE FILE) OR IF THERE IS NO SUCH SOFTWARE LICENSE AGREEMENT OR CLICKWRAP END USER LICENSE AGREEMENT, THE LICENSE(S) LOCATED IN THE "LICENSE" FILE(S) OF THE SOFTWARE. USE OF THIS DOCUMENT IS SUBJECT TO THOSE TERMS AND CONDITIONS, AND YOUR USE HEREOF SHALL CONSTITUTE ACCEPTANCE OF AND AN AGREEMENT TO BE BOUND BY THE SAME.

This document is subject to U.S. and international copyright laws and treaties. No part of this document may be reproduced in any form without the written authorization of TIBCO Software Inc.

TIBCO, the TIBCO logo, the TIBCO O logo, TIBCO Hawk, TIBCO Rendezvous, TIBCO Runtime Agent, ActiveMatrix BusinessWorks, TIBCO Administrator, TIBCO Designer, TIBCO ActiveMatrix Service Gateway, TIBCO BusinessEvents, TIBCO BusinessConnect, and TIBCO BusinessConnect Trading Community Management are either registered trademarks or trademarks of TIBCO Software Inc. in the United States and/or other countries.

Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle Corporation and/or its affiliates.

All other product and company names and marks mentioned in this document are the property of their respective owners and are mentioned for identification purposes only.

This software may be available on multiple operating systems. However, not all operating system platforms for a specific software version are released at the same time. See the readme file for the availability of this software version on a specific operating system platform.

THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS DOCUMENT COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THIS DOCUMENT. TIBCO SOFTWARE INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS DOCUMENT AT ANY TIME.

THE CONTENTS OF THIS DOCUMENT MAY BE MODIFIED AND/OR QUALIFIED, DIRECTLY OR INDIRECTLY, BY OTHER DOCUMENTATION WHICH ACCOMPANIES THIS SOFTWARE, INCLUDING BUT NOT LIMITED TO ANY RELEASE NOTES AND "READ ME" FILES.

This and other products of TIBCO Software Inc. may be covered by registered patents. Please refer to TIBCO's Virtual Patent Marking document (<https://www.tibco.com/patents>) for details.

Copyright © 1999-2021. TIBCO Software Inc. All Rights Reserved.