

TIBCO Runtime Agent™

Installing Into a Cluster

Software Release 5.10

August 2015

Important Information

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 contains confidential information that 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, Two-Second Advantage, TIBCO Hawk, TIBCO Rendezvous, TIBCO Runtime Agent, TIBCO 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.

Enterprise Java Beans (EJB), Java Platform Enterprise Edition (Java EE), Java 2 Platform Enterprise Edition (J2EE), and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle Corporation in the U.S. and other countries.

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 Product is covered by U.S. Patent No. 6,970,981.

Copyright © 1998-2015 TIBCO Software Inc. ALL RIGHTS RESERVED.

TIBCO Software Inc. Confidential Information

Contents

Figures	v
Tables	vii
Preface	ix
Changes from the previous Release of this Guide	x
Related Documentation	xi
TIBCO Runtime Agent Documentation	xi
Other TIBCO Product Documentation	xi
Typographical Conventions	xiii
Connecting with TIBCO Resources	xvi
How to Join TIBCOCommunity	xvi
How to Access TIBCO Documentation	xvi
How to Contact TIBCO Support	xvi
Chapter 1 Concepts	1
Overview	2
Cluster Aware and Cluster Unaware Applications	2
High-availability Modes	3
Basic Cluster Terminology	3
Failover Cluster Characteristics	5
Cluster Support Limitations	7
Placing TIBCO Applications into a Cluster Scenarios	8
Entire Domain in a Single Cluster Group	9
Entire Domain Split Into Multiple Cluster Groups	10
Domain Not in Cluster and Services in Cluster	11
Chapter 2 Installing the Software Into a Cluster	13
Overview	14
Example Installation Diagram	15
Creating a Cluster Group	17
Installing TIBCO Runtime Agent on the Active Node	18
Installing TIBCO Administrator on the Active Node	19
Create the Initial Domain	19

- Register the Services with the Cluster 24
 - Add Service Name to the Cluster 24
 - Change a Service Property 28
- Configure the Second Node 30
 - Install TIBCO Software 31
 - Join the Second Node to the Domain 31
 - Bring the Service Online 33
- Adding a TIBCO Application as a Service to the Cluster 34
 - Add the Process as a Service 34
 - Change a Service Property 39
 - Bring the Service Online 40
- Adding a TIBCO Application as an Application to a Cluster 41
 - Add the Process as an Application 41
 - Change an Application Property 44
 - Bring the Application Online 45
- Installing TIBCO Software Onto a Shared Drive 46
- Verify the Installation 47
- Installed Files 49
- Chapter 3 Deploying and Starting Applications Under Cluster Control 51**
 - Redeploying Applications 52
 - Forcing Redeployment of an NT Service 53
 - Starting and Stopping Applications Managed by Cluster Software 54
 - Removing Orphaned Services 55
- Index 57**

Figures

Figure 1	Failover Cluster Solution	2
Figure 2	Failover Cluster Dependency Order.	6
Figure 3	Applications in a Domain Configured in a Cluster Group.	9
Figure 4	Dividing the Application's Service Instances into Different Cluster Groups	10
Figure 5	Domain Not in Cluster and Services in Cluster	12
Figure 6	Example Installation Cluster.	15
Figure 7	Cluster Administrator Window	17
Figure 8	Enter Domain Details	19
Figure 9	Change the Node Name	20
Figure 10	Select Show Advanced Check Box	20
Figure 11	Domain Home Paths Configuration	21
Figure 12	Virtual IP Address Field	22
Figure 13	DomainUtility Confirm Window.	23
Figure 14	Register the Services with the Cluster	24
Figure 15	Cluster Administartor Window	25
Figure 16	New Resource Dialog	25
Figure 17	Select Source Type	26
Figure 18	Possible Owners Dialog	26
Figure 19	Dependencies Dialog	27
Figure 20	Generic Service Parameters	27
Figure 21	Registry Replication	28
Figure 22	Cluster Administrator Window	28
Figure 23	Advanced Tab	29
Figure 24	Right-click the Group and Select Move	30
Figure 25	Active Second Move	30
Figure 26	Select Server	31
Figure 27	Select a Task	32
Figure 28	Configuration Dialog.	32

Figure 29	Services Console	33
Figure 30	Bring Online.....	33
Figure 31	Service Name Field.....	35
Figure 32	Cluster Administrator Window.....	35
Figure 33	Provide Name for Resource	36
Figure 34	Resource Type Field.....	36
Figure 35	Possible Owners Window	37
Figure 36	Dependencies Dialog	37
Figure 37	Enter Service Name	38
Figure 38	Registry Replication Dialog.....	38
Figure 39	Correct Registration of Resource	39
Figure 40	Cluster Administrator.....	39
Figure 41	BW Service Properties Window.....	40
Figure 42	Bringing Service Online.....	40
Figure 43	Application's .cmd File Location.....	42
Figure 44	Cluster Administrator Window	42
Figure 45	Select Resource Type as Generic Application	43
Figure 46	Dependencies Dialog	43
Figure 47	Generic Application Parameter Window	44
Figure 48	Correct ResourceRegistration.....	44
Figure 49	To Bring the Application Online.....	45
Figure 50	TIBCO Administrator GUI Login Screen	47
Figure 51	Virtual Machine	48
Figure 52	Using Cluster Administrator Software to Start and Stop the Application.....	52
Figure 53	Force Redeployment.....	53

Tables

Table 1 General Typographical Conventions xiii

Table 2 Syntax Typographical Conventions xiv

Preface

This document gives comprehensive instructions on how to install TIBCO Runtime Agent™ software and deploy applications in a cluster environment.

Topics

- [Changes from the previous Release of this Guide, page x](#)
- [Related Documentation, page xi](#)
- [Typographical Conventions, page xiii](#)
- [Connecting with TIBCO Resources, page xvi](#)

Changes from the previous Release of this Guide

All the screenshots have been updated with new TIBCO logo.

Related Documentation

This section lists documentation resources you may find useful.

TIBCO Runtime Agent Documentation

The TIBCO Runtime Agent™ software suite is a prerequisite for other TIBCO software products. In addition to Runtime Agent components, the software suite includes the third-party libraries used by other TIBCO products, TIBCO Designer™, Java Runtime Environment (JRE), TIBCO Rendezvous®, and TIBCO Hawk®.

The following documents form the TIBCO Runtime Agent™ documentation set:

- *TIBCO Runtime Agent™ Installation* Read this manual for instructions on site preparation and installation.
- *TIBCO Runtime Agent™ Installing Into a Cluster* Read this manual for instructions on installing TIBCO applications into a cluster environment.
- *TIBCO Runtime Agent™ Upgrading to Release 5.10.0* Read this manual for instructions on upgrading from release 5.x to release 5.10.0
- *TIBCO Runtime Agent™ Domain Utility User's Guide* Read this manual for instructions on using TIBCO Domain Utility to create and manage administration domains.
- *TIBCO Runtime Agent™ Scripting Deployment User's Guide* Read this manual for instructions on using the AppManage scripting utility to deploy applications.
- *TIBCO Runtime Agent™ Authentication API User's Guide* Read this manual for instructions on using Authentication API.
- *TIBCO Runtime Agent™ Release Notes* Read the release notes for a list of new and changed features. This document also contains lists of known issues and closed issues for this release.

Other TIBCO Product Documentation

You may find it useful to read the documentation for the following TIBCO products:

- **TIBCO Administrator™** : TIBCO Administrator allows you to manage users, machines and applications defined in a TIBCO administration domain. The TIBCO Administrator graphical user interface enables users to deploy, monitor, and start and stop TIBCO applications.

- TIBCO Designer™: This graphical user interface is used for designing and creating integration project configurations and building an Enterprise Archive (EAR) for the project. The EAR can then be used by TIBCO Administrator for deploying and running the application.
- TIBCO Hawk®: This is a tool for monitoring and managing distributed applications and operating systems.
- TIBCO Rendezvous®: Rendezvous enables programs running on many different kinds of computers on a network to communicate seamlessly. It includes two main components: the Rendezvous application programming interface (API) in several languages, and the Rendezvous daemon.
- TIBCO Enterprise Message Service™: This software lets application programs send and receive messages using the Java Message Service (JMS) protocol. It also integrates with TIBCO Rendezvous and TIBCO SmartSockets® messaging products.
- TIBCO ActiveMatrix BusinessWorks™: ActiveMatrix BusinessWorks is a scalable, extensible, and easy to use integration platform that allows you to develop integration projects. ActiveMatrix BusinessWorks includes a GUI for defining business processes and an engine that executes the process.
- TIBCO® Adapter software: TIBCO Runtime Agent is a prerequisite for TIBCO Adapter products. You will therefore find TIBCO Adapter product documentation useful.

Typographical Conventions

The following typographical conventions are used in this manual.

Table 1 General Typographical Conventions

Convention	Use
<i>ENV_NAME</i> <i>TIBCO_HOME</i> <i>TRA_HOME</i>	<p>TIBCO products are installed into an installation environment. A product installed into an installation environment does not access components in other installation environments. Incompatible products and multiple instances of the same product must be installed into different installation environments.</p> <p>An installation environment consists of the following properties:</p> <ul style="list-style-type: none"> • Name Identifies the installation environment. This name is referenced in documentation as <i>ENV_NAME</i>. On Microsoft Windows, the name is appended to the name of Windows services created by the installer and is a component of the path to the product shortcut in the Windows Start > All Programs menu. • Path The folder into which the product is installed. This folder is referenced in documentation as <i>TIBCO_HOME</i>. <p><i>TIBCO Runtime Agent</i> installs into a directory within a <i>TIBCO_HOME</i>. This directory is referenced in documentation as <ProductAcronym>_HOME. The default value of <ProductAcronym>_HOME depends on the operating system. For example on Windows systems, the default value is C:\tibco\<ProductAcronym>\<ReleaseNumber>.</p>
code font	<p>Code font identifies commands, code examples, filenames, pathnames, and output displayed in a command window. For example:</p> <p>Use MyCommand to start the foo process.</p>
bold code font	<p>Bold code font is used in the following ways:</p> <ul style="list-style-type: none"> • In procedures, to indicate what a user types. For example: Type admin. • In large code samples, to indicate the parts of the sample that are of particular interest. • In command syntax, to indicate the default parameter for a command. For example, if no parameter is specified, MyCommand is enabled: MyCommand [enable disable]

Table 1 General Typographical Conventions (Cont'd)




Convention	Use
<i>italic font</i>	<p>Italic font is used in the following ways:</p> <ul style="list-style-type: none">• To indicate a document title. For example: See <i>TIBCO ActiveMatrix BusinessWorks Concepts</i>.• To introduce new terms For example: A portal page may contain several portlets. <i>Portlets</i> are mini-applications that run in a portal.• To indicate a variable in a command or code syntax that you must replace. For example: <code>MyCommand <i>PathName</i></code>
Key combinations	<p>Key name separated by a plus sign indicate keys pressed simultaneously. For example: <code>Ctrl+C</code>.</p> <p>Key names separated by a comma and space indicate keys pressed one after the other. For example: <code>Esc, Ctrl+Q</code>.</p>
	<p>The note icon indicates information that is of special interest or importance, for example, an additional action required only in certain circumstances.</p>
	<p>The tip icon indicates an idea that could be useful, for example, a way to apply the information provided in the current section to achieve a specific result.</p>
	<p>The warning icon indicates the potential for a damaging situation, for example, data loss or corruption if certain steps are taken or not taken.</p>

Table 2 Syntax Typographical Conventions

Convention	Use
[]	<p>An optional item in a command or code syntax.</p> <p>For example:</p> <pre>MyCommand [optional_parameter] required_parameter</pre>
	<p>A logical OR that separates multiple items of which only one may be chosen.</p> <p>For example, you can select only one of the following parameters:</p> <pre>MyCommand param1 param2 param3</pre>

Table 2 Syntax Typographical Conventions (Cont'd)

Convention	Use
{ }	<p>A logical group of items in a command. Other syntax notations may appear within each logical group.</p> <p>For example, the following command requires two parameters, which can be either the pair param1 and param2, or the pair param3 and param4.</p> <pre>MyCommand {param1 param2} {param3 param4}</pre> <p>In the next example, the command requires two parameters. The first parameter can be either param1 or param2 and the second can be either param3 or param4:</p> <pre>MyCommand {param1 param2} {param3 param4}</pre> <p>In the next example, the command can accept either two or three parameters. The first parameter must be param1. You can optionally include param2 as the second parameter. And the last parameter is either param3 or param4.</p> <pre>MyCommand param1 [param2] {param3 param4}</pre>

Connecting with TIBCO Resources

How to Join TIBCOCommunity

TIBCOCommunity is an online destination for TIBCO customers, partners, and resident experts. It is a place to share and access the collective experience of the TIBCO community. TIBCOCommunity offers forums, blogs, and access to a variety of resources. To register, go to <http://www.tibcommunity.com>.

How to Access TIBCO Documentation

You can access TIBCO documentation here:

<http://docs.tibco.com>

How to Contact TIBCO Support

For comments or problems with this manual or the software it addresses, contact TIBCO Support as follows:

- For an overview of TIBCO Support, and information about getting started with TIBCO Support, visit this site:

<http://www.tibco.com/services/support>

- If you already have a valid maintenance or support contract, visit this site:

<https://support.tibco.com>

Entry to this site requires a user name and password. If you do not have a user name, you can request one.

Chapter 1 **Concepts**

This chapter explains concepts you should be familiar with before installing TIBCO applications in a cluster.

Topics

- [Overview, page 2](#)
- [Failover Cluster Characteristics, page 5](#)
- [Cluster Support Limitations, page 7](#)
- [Placing TIBCO Applications into a Cluster Scenarios, page 8](#)

Overview

A cluster is a group of machines that work together as a single system to ensure that applications and resources are available at all times. The machines are managed as a single system using cluster software, which provides a way to support fault-tolerance, high-availability, scalability, and so on.


The computers are physically connected by cables and programmatically connected by cluster software. These connections allow computers to use failover and load balancing, which is not possible with a stand-alone computer.

Figure 1 shows a failover cluster solution. The nodes in a cluster are configured with a shared drive. The cluster software manages access to the shared drive so that only applications on a single machine use a particular resource at anytime.

The failover cluster solution primarily provides uninterrupted service in the event of a failure within the cluster. This solution is commonly used for database applications.

Applications that are deployed to a failover cluster can be managed under cluster software. Applications are under the control of cluster software, which ensures that only one application runs at once on one of the cluster machines. This kind of application is in active-passive mode.

Figure 1 Failover Cluster Solution



Cluster Aware and Cluster Unaware Applications

A cluster environment supports two types of applications: cluster-aware applications and cluster-unaware applications.



TIBCO applications are cluster-unaware. The applications can be run seamlessly in a cluster environment, but do not use any cluster features.

- Cluster-aware applications run in the cluster and are aware if the cluster itself is running, or whether a resource needed by the application is available. A cluster-aware application is supported by a resource type that is assigned by the cluster software.
- Cluster-unaware applications know nothing about the cluster environment and consequently cannot exchange data with cluster objects or be aware that they are running under cluster software control.

Cluster software manages cluster-unaware applications by using basic methods for failure detection and application shutdown. For example, the Microsoft Cluster Server solution manages cluster-unaware software as a generic application resource type.

High-availability Modes

High-availability applications can be categorized to different modes, hot standby, warm standby, and cold standby in a cluster.



When TIBCO applications are installed under a cluster, high-availability is best characterized as warm-standby.

- Hot standby—Applications on the backup node take over with almost no downtime and may even approach zero downtime. Using hot standby, two processors use hardware checkpoints to verify synchronization after each CPU instruction.
- Warm standby—A slight delay occurs when a backup application takes over for another application.
- Cold standby—The delay is a little longer than warm standby. In cold standby, the cluster software can't detect an outage and applications must be restarted manually. One or more spare processors are set aside to be used if an in-service processor fails.

Basic Cluster Terminology

The following basic terminology is used in this document.

- Node. The term used to refer to a server that is a member of a cluster.
- Resource. A hardware or software component that exists in a cluster, such as a disk, an IP address, a network name or an application.
- Group. A combination of resources that are managed as a unit of failover. Groups are also known as resource groups, cluster packages, or service groups.

- **Dependency.** An alliance between two or more resources in the cluster architecture.
- **Failover or failback.** The process of moving resources from one server to another. Failover can occur when one server experiences a failure of some sort or when an administrator initiates a proactive failover.
- **Active/Active.** From a software perspective, this describes applications (or resources) that can exist as multiple instances in a cluster. This means that both nodes can be active servicing clients.
- **Active/Passive.** This term describes applications that run as a single instance in a cluster. In general, this also means that one node typically is idle until a failover occurs.
- **Rolling Upgrade.** A rolling upgrade allows you to upgrade the software on cluster nodes, one node at a time, so that the applications and resources on the cluster are always available even though the nodes being upgraded are not.
- **Shared storage.** This refers to the external SCSI or fibre channel storage enclosure and the disks contained therein. Shared storage is a requirement for multi-node clusters. Although this storage is shared, only one node can access an external storage resource at any given time.

Failover Cluster Characteristics

Each failover cluster may be divided into multiple failover cluster groups, sometimes called cluster packages or failover service groups. Because each cluster group must completely own its resources and a shared disk is a cluster resource type, the maximum number of cluster groups in a cluster are determined by the number of shared cluster disks that a cluster contains.

A failover cluster provides high availability and fault-tolerance, but not load balancing as each failover cluster group must run from only one node at a time.

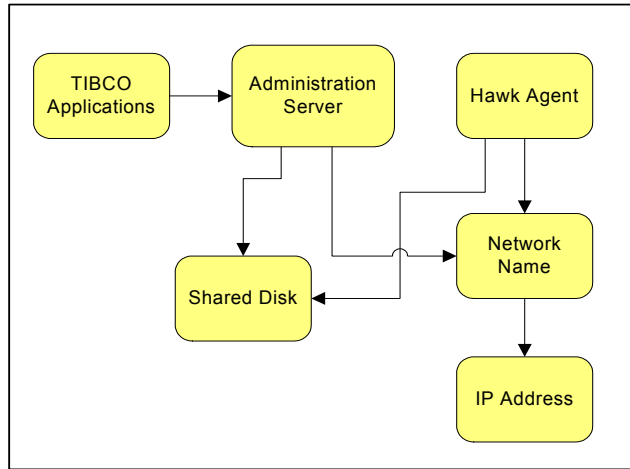
Applications deployed in a failover cluster are under the control of the cluster administrator and run in active-passive mode. For the same application, only one service instance runs in a cluster group at a time. Since a cluster group can belong to only one cluster node at any point of time, the node currently owning the group hosts the application.

TIBCO applications can be deployed to a failover cluster under one or multiple cluster groups. The cluster group that hosts TIBCO products must own the network name (optional), IP address and shared disk resources. If a cluster group lacks any of these required resources, it cannot host TIBCO products correctly as the products depend on the availability of those resources. TIBCO applications are resources in the group that depend on these basic resources.

Applications deployed in a cluster group are managed by cluster software as one atomic unit. If a failover occurs, the entire group must be moved to another node in the cluster group.

Applications with interdependencies must be placed in the same cluster group. The dependencies must be configured on the resource group so that the cluster service can control the order in which resources are brought online and offline.

For example, [Figure 2](#) shows the dependency order in which the cluster brings the resources back online after a failover. An arrow from an object points to its dependent object. If the network name resource is not defined, the IP address is used instead.

Figure 2 Failover Cluster Dependency Order

If your domain is configured to use local application data, a local DAT file is used at runtime instead of server-based DAT file and applications can be started even if the administration server is not yet online. In this case there would be no dependency between the administration server and TIBCO applications.

Cluster Support Limitations

Some cluster vendors, such as Veritas allow the use of parallel groups. A parallel group can run concurrently on multiple systems in the cluster at a time. A parallel service group is more complex than a failover group. It requires an application that can safely be started on more than one system at a time with no threat of data corruption, or that the data being accessed is local to each server.



TIBCO Administrator does not support parallel groups. The main architecture limitation is that only one copy of a TIBCO Hawk Agent instance can be running from one system in the cluster at a time.

Cluster vendors also provide load balanced cluster groups. A load balanced cluster solution helps to implement an enterprise-wide scalable solution for incoming TCP/IP traffic. This solution is commonly used for web server applications.

Applications that are deployed to a load balanced cluster run separate copies of the applications on each host machine. Network load balancing distributes incoming client requests across the machines in the cluster.



TIBCO applications are intended to work in failover cluster environments, not in load balanced clusters. Load-balanced cluster solutions are not addressed in this document.

Placing TIBCO Applications into a Cluster Scenarios

This section presents scenarios that show how to place TIBCO applications into a cluster. You should first decide how to organize an administration domain into cluster groups. You can place your entire domain within a same cluster group, or divide your domain into multiple cluster groups. You can also keep some of the domain software outside the cluster and place only certain applications under cluster software control.

The trade-off is often manageability versus finer failover control and better computer resource usage. You should also consider how to provide high availability of external products that TIBCO products rely on, such as databases.

Note that each cluster group:

- must have its own shared drive. A cluster group's shared drive cannot be shared by another cluster group.
- must have a single virtual IP address or network name that maps to the virtual IP address.

When you configure TIBCO applications to work with a cluster solution, the applications (TIBCO Runtime Agent, TIBCO Administrator, TIBCO BusinessWorks, and so on) are typically installed on each cluster group node's local drive. The TIBCO home location must be the same on each node in the cluster group.

When you use TIBCO Domain Utility to create an administration domain for the cluster group, the TIBCO Runtime Agent domain home and the TIBCO Administrator domain home is set to the cluster group's shared drive. The cluster group's virtual network name or IP address is also specified.

When you deploy applications from a node in a cluster group, deployment files are written to the cluster group's shared drive. If the domain was created with the Local Application Data option enabled, all files required for runtime are written to the shared drive.




Applications that are managed by a cluster should not be started or stopped using the TIBCO Administrator GUI. If you stop an application from TIBCO Administrator, the cluster would assume the application had failed. If the application is configured to be restarted by the cluster when it fails, it will be automatically restarted by the cluster. See [Deploying and Starting Applications Under Cluster Control](#) on page 51.

Entire Domain in a Single Cluster Group

You can group TIBCO applications along with their related resources such as IP address, network name and shared disk into just one cluster group. One cluster group corresponds to one logical machine as defined in the administration domain.

Figure 3 shows all applications in a domain configured in a cluster group. The domain includes the TIBCO Hawk Agent, TIBCO Administrator administration server, and all managed applications in the domain. The Hawk Agent and administration server must be controlled by the cluster software. Other TIBCO applications such as the TIBCO Enterprise Message Service server could also be configured into the cluster group.

Figure 3 Applications in a Domain Configured in a Cluster Group



Scenario Benefits

A domain in one cluster group is easier to manage. Placing the entire domain in one cluster group results in just one virtual machine in the TIBCO Administrator domain. The name of the virtual machine is the network name resource of the cluster group, while the IP address of the machine is the IP address resource of the cluster group. If a network name is not configured, the virtual IP address is used as the machine name.

Scenario Risks

An entire domain in a single cluster group is easier to manage but, because a cluster group is an atomic unit for failover control, the entire domain will failover all together if one resource in the group fails.

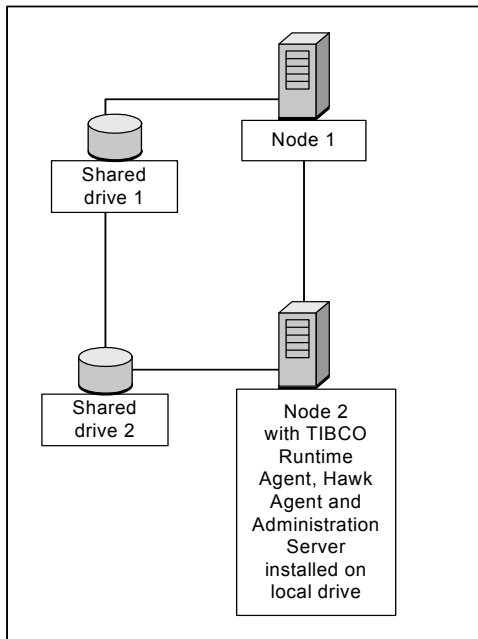
Also, because a single cluster group is not optimized to use computer resources, all applications in a cluster group must be running from one cluster node at any given time. This means the entire TIBCO domain runs on the active cluster node, while the other cluster nodes are idle, waiting for failover.

Entire Domain Split Into Multiple Cluster Groups

A cluster group failover occurs as a single atomic group. If you do not want TIBCO Administrator and its managed applications to failover as an atomic group, you can divide the application's service instances into different cluster groups as shown in the next diagram. You can also divide the managed applications in the domain into multiple cluster groups to get a more granular control over failover.

In this scenario, two cluster groups are possible, one with Node 1, Node 2 and shared drive 1, and another with Node 1, Node 2 and shared drive 2.

Figure 4 Dividing the Application's Service Instances into Different Cluster Groups



If you decide to divide a domain into multiple cluster groups, make sure each cluster group owns its own shared disk, network name (optional) and IP address. In a TIBCO Administrator domain, you will see multiple virtual machines. Each machine represents a cluster group. The name of each machine is the network name resource of a cluster group, while the IP address of the machine is the IP address resource of the cluster group.

The cluster software ensures that at any one point of time the shared disk can only be accessed by the owning node of the cluster group. Other cluster nodes cannot access the shared disk even though all nodes in the cluster groups are physically connected to the shared disk.

Scenario Benefits

A domain divided into multiple cluster groups provides more granular control of failover and more optimized use of computer resources in a cluster environment.

For example, if you have two services A and B running in a cluster with one cluster group, both A and B must run in the same cluster group and the group must run either on machine 1 or machine 2. One machine hosts both services A and B while the other machine is idle.

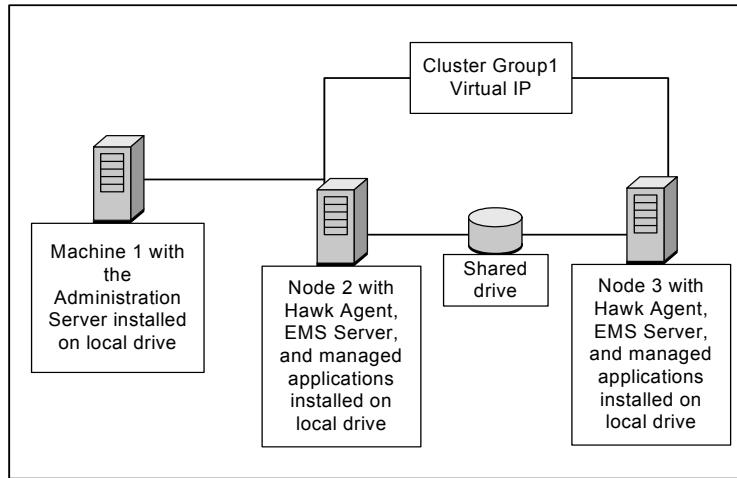
If you form two cluster groups, service A can be placed in the cluster group that runs on machine 1. Service B can be placed in the cluster group 2 that runs on machine B. This allows each machine to run a service, and be a backup for the other machine. No machine is idle.

Scenario Risks

A domain divided into multiple cluster groups can be more difficult to manage. For example, not only do the resources within each cluster group have order dependencies, the cluster groups also have order dependencies. The cluster group that hosts TIBCO Administrator must be online before applications in other cluster groups that depend on TIBCO Administrator can be brought online.

Domain Not in Cluster and Services in Cluster

[Figure 5](#) shows a scenario where the TIBCO Administrator administration server is installed on a machine that is not under cluster control. Other TIBCO applications such as the TIBCO Hawk Agent, TIBCO Enterprise Message Service server and applications managed in the domain are configured in a cluster group.

Figure 5 Domain Not in Cluster and Services in Cluster

Scenario Benefits

Only the applications that need to be managed by the cluster are under cluster control. This allows for better use of cluster resources in the event of a failover.

Scenario Risks

Applications not managed by the cluster do not have failover control. However, primary and secondary administration servers can be configured for fault tolerance.

Chapter 2

Installing the Software Into a Cluster

This chapter explains how to install and configure TIBCO applications in a cluster environment.

Topics

- [Overview, page 14](#)
- [Creating a Cluster Group, page 17](#)
- [Installing TIBCO Runtime Agent on the Active Node, page 18](#)
- [Installing TIBCO Administrator on the Active Node, page 19](#)
- [Register the Services with the Cluster, page 24](#)
- [Configure the Second Node, page 30](#)
- [Adding a TIBCO Application as a Service to the Cluster, page 34](#)
- [Verify the Installation, page 47](#)
- [Installed Files, page 49](#)

Overview

TIBCO Runtime Agent, TIBCO Administrator, and the applications it manages need not be modified to run in a cluster environment. You can install the software on a cluster group node's local drive, the cluster group's shared drive, or a mapped or mounted network drive. The following points apply to each installation:

- Login to each cluster node as the same user to perform the installation.
- Each TIBCO application must be registered directly with the cluster software. This is done using your cluster software.
- Applications running in a cluster group that are managed by TIBCO Administrator must be in warm-standby mode, not in hot-standby mode and must all fail-over together as an atomic unit. This is specified in your cluster software.
- If a persistent TCP connection is maintained by the application, each client application must re-establish its connection to the administration server after a cluster failover. This is specified in your cluster software.



See the *TIBCO Runtime Agent Release Notes* for known issues about clusters that affect TIBCO applications.

A primary administration server that is configured in a cluster environment typically does not have a secondary server configured because the cluster provides a backup. However, you can configure one or multiple secondary servers when the primary server is under cluster software control. The secondary servers are not required to be configured on a standby node in the cluster.

When monitoring applications in warm-standby mode under the control of cluster software, applications are deployed to a virtual machine and not to physical machines. At any point of time, the deployed application can only be running on one node in the cluster. The TIBCO Hawk service and the applications it monitors must belong to the same cluster resource group.

Before you use TIBCO Domain Utility to configure the administration domain, you should decide how many cluster groups your domain will be divided among and form the cluster groups.

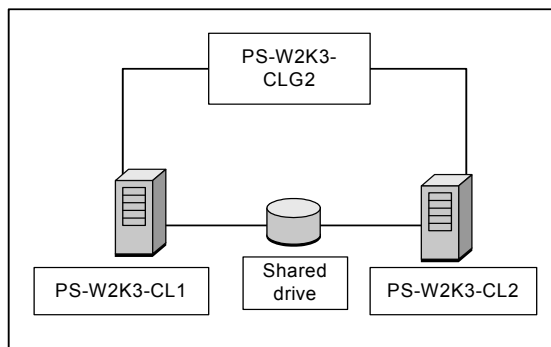
It is recommended that TIBCO Runtime Agent, TIBCO Administrator and other TIBCO applications be installed on each cluster node's local disk, and not on the cluster group's shared disk. Installing TIBCO products on a local disk rather than a cluster group's shared disk enables rolling upgrades, and allows TIBCO products to be placed in multiple cluster groups.

Example Installation Diagram

This section explains how to install TIBCO products in a cluster environment. This example uses Microsoft Cluster Server software.

Figure 6 shows the cluster that is used in this installation example. The virtual cluster name is PS-W2K3-CLG2. The cluster has two nodes, PS-W2K3-CL1 and PS-W2K3-CL2.

Figure 6 Example Installation Cluster



TIBCO applications will be installed on each node's local drive, and then registered with the cluster software. Additionally an administration domain will be created on the shared drive.

PS-W2K3-CL1 Node

The following will be done on this node:

- Install TIBCO Runtime Agent and TIBCO Administrator. Other TIBCO applications, such as TIBCO BusinessWorks can be installed. This is not shown in the example.
- Using your cluster administrator software, make sure that PS-W2K3-CL1 is the active node.
- Create an administration domain on the shared drive. This is done once, after TIBCO applications are installed on the first node.
- Register the administration server service and the TIBCO Hawk Agent service with the cluster software. This is done once. It need not be done when other nodes are configured.

PS-W2K3-CL2 Node

The following will be done on this node:

- Install TIBCO Runtime Agent and TIBCO Administrator.
- Other TIBCO applications, such as TIBCO BusinessWorks can be installed.
- Using your cluster administrator software, make sure that PS-W2K3-CL1 is the active node.
- Using TIBCO Domain Utility, add the logical machine to the administration domain.

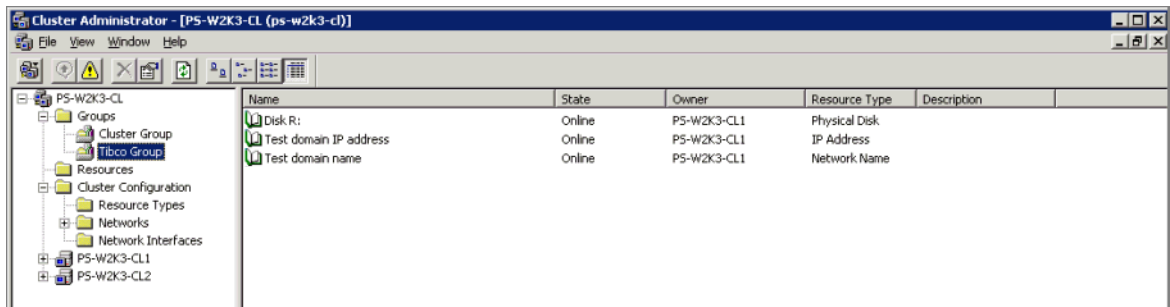
Creating a Cluster Group

Before starting the installation, you must setup the cluster software by creating a cluster group and assigning it the following resources:

- Physical Disk
- IP Address
- Network Name (optional)

Figure 7 shows the cluster administrator window for the PS-W2K3-CL cluster. The cluster includes a group named Tibco-group that includes the above resources. Although not shown in the diagram, the cluster virtual hostname has been defined as PS-W2K3-CLG2. Two nodes are defined in the cluster, PS-W2K3-CL1 and PS-W2K3-CL2. The dialog indicates that the current owner of each resource is node PS-W2K3-CL1. The TIBCO software will be installed on this node first.

Figure 7 Cluster Administrator Window



Installing TIBCO Runtime Agent on the Active Node

Install TIBCO Runtime Agent 5.7 or greater on the local drive attached to the active node (see [Installing TIBCO Software Onto a Shared Drive on page 46](#) for detailed instructions). For information about installing the software, see *TIBCO Runtime Agent Installation Guide*.

The installation home must be identical on each node in the cluster. The installation home is set when installing TIBCO Runtime Agent. In our example, because TIBCO Runtime Agent is installed on PS-W2K3-CL1 in `c:/tibco`, TIBCO Runtime Agent must be installed on PS-W2K3-CL2 in `c:/tibco`. The same version of TIBCO Runtime Agent software must be installed on each node in the cluster group.

Installing TIBCO Administrator on the Active Node

Install TIBCO Administrator 5.7 or greater on the local drive attached to the active node (see [Installing TIBCO Software Onto a Shared Drive on page 46](#) for detailed instructions). For information about installing the software, see the *TIBCO Administrator Installation Guide*.

Create the Initial Domain

After TIBCO Administrator is installed, TIBCO Domain Utility launches so the initial administration domain can be created. Because the domain will be created on the shared drive and will use the cluster virtual host name, you must specify the machine name, administration server home and application domain home.



This step is done only once when TIBCO Administrator is installed on the first active node. If you are installing TIBCO Administrator on the second node, do not complete this section and go to [Join the Second Node to the Domain on page 31](#).

1. In Domain Utility, click next until the following screen appears. Provide a name for the domain in the project directory in the Administration Domain field.

Figure 8 Enter Domain Details

TIBCO DomainUtility 5.9 - Administrator Configuration

Please provide appropriate values for the following fields.

Domain Details

Administration Domain: ClusterDemo

Project Directory: \\nistrator\\domain\\ClusterDemo\\data

Machine: PS-W2K3-CL1

Hawk Cluster: 10.97.97.0

Encoding: ISO8859-1

Domain Data Store Configuration

☐ User and Group information retrieved from a corporate LDAP

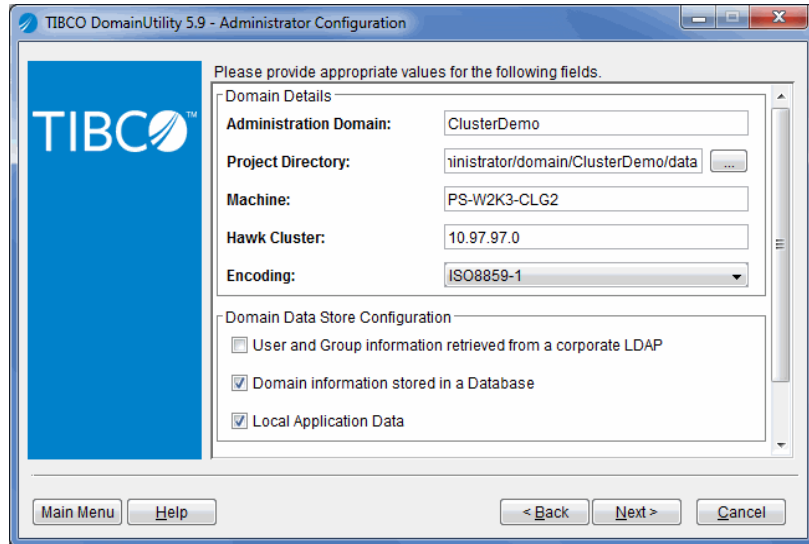
☒ Domain information stored in a Database

☒ Local Application Data

Main Menu Help < Back Next > Cancel

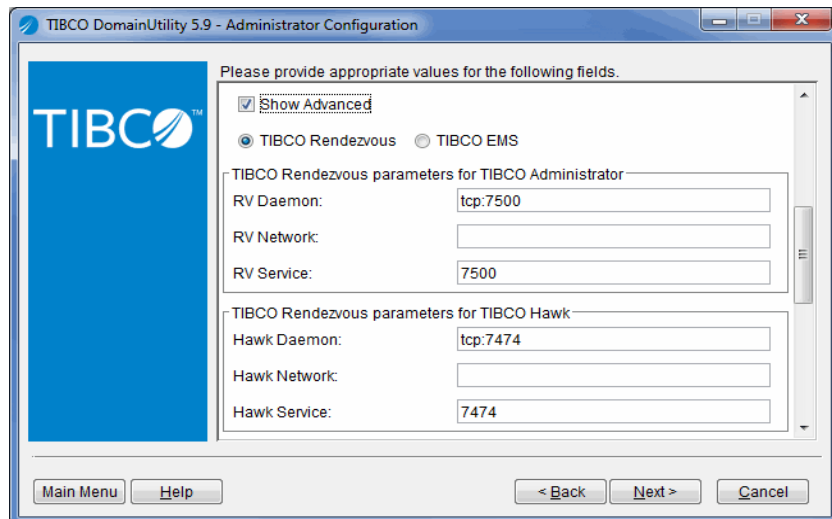
2. In the Machine field, change the given name to the cluster virtual hostname. In this example, the node name PS-W2K3-CL1 is changed to the cluster virtual hostname, PS-W2K3-CLG2.

Figure 9 Change the Node Name



3. Click the Show Advanced check box.

Figure 10 Select Show Advanced Check Box



4. Scroll down to the Domain Home Paths Configuration section. Change the fields to the location of the shared drive. For example, the next diagram shows the fields that contains the location of the TIBCO Runtime Agent home and the administration domain home.

Figure 11 Domain Home Paths Configuration

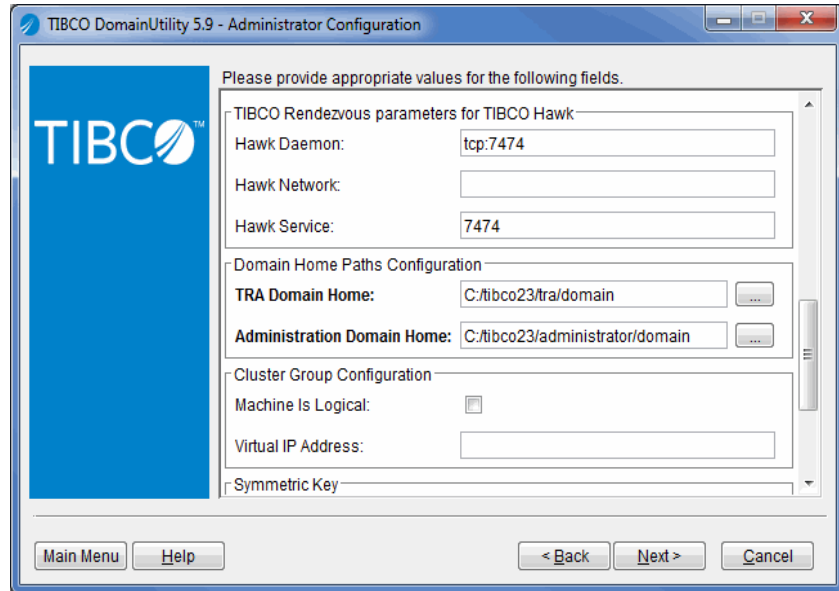
The screenshot shows the 'TIBCO DomainUtility 5.9 - Administrator Configuration' window. It features a blue sidebar with the TIBCO logo. The main content area is titled 'Please provide appropriate values for the following fields.' and contains three sections:

- TIBCO Rendezvous parameters for TIBCO Hawk:**
 - Hawk Daemon: tcp:7474
 - Hawk Network: (empty field)
 - Hawk Service: 7474
- Domain Home Paths Configuration:**
 - TRA Domain Home: C:/tibco23/tra/domain
 - Administration Domain Home: C:/tibco23/administrator/domain
- Cluster Group Configuration:**
 - Machine Is Logical: ☒
 - Virtual IP Address: (empty field)
 - Symmetric Key: (empty field)

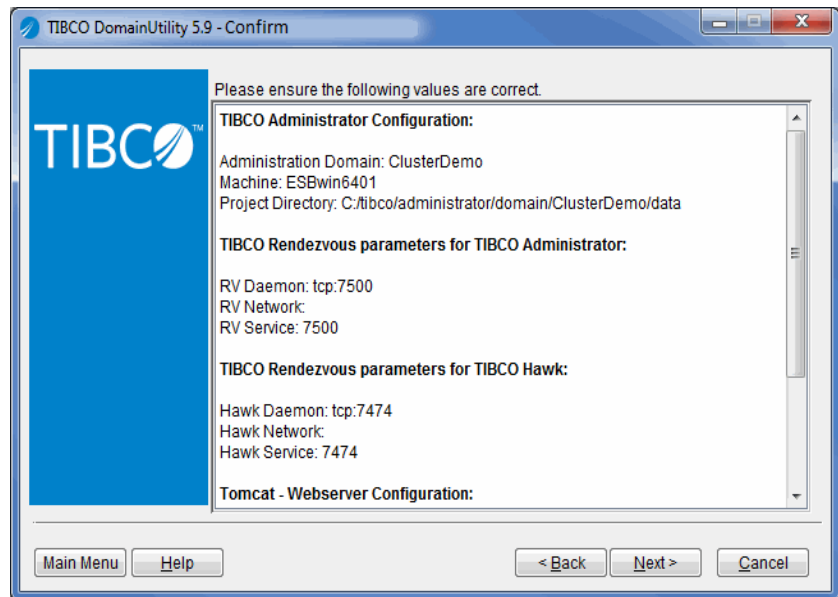
At the bottom, there are buttons for 'Main Menu', 'Help', '< Back', 'Next >', and 'Cancel'.

5. The Machine is Logical check box must be selected if the value entered in the Machine field (at the top) is a logical machine name. The Virtual IP Address field is optional and can contain the IP address assigned to the virtual cluster host.

Figure 12 Virtual IP Address Field



6. Click Next. In most cases the default Tomcat server web ports are used.
7. Click Next. Provide the name and password for the domain administrator user.
8. Click Next. Review the options you have selected. You can click Back to change values, if needed.
9. Click Next. TIBCO Domain Utility creates the domain and lists the NT Services that should be started. Note that you should not start the services from the Services panel. Instead, the services must be registered with the cluster software.

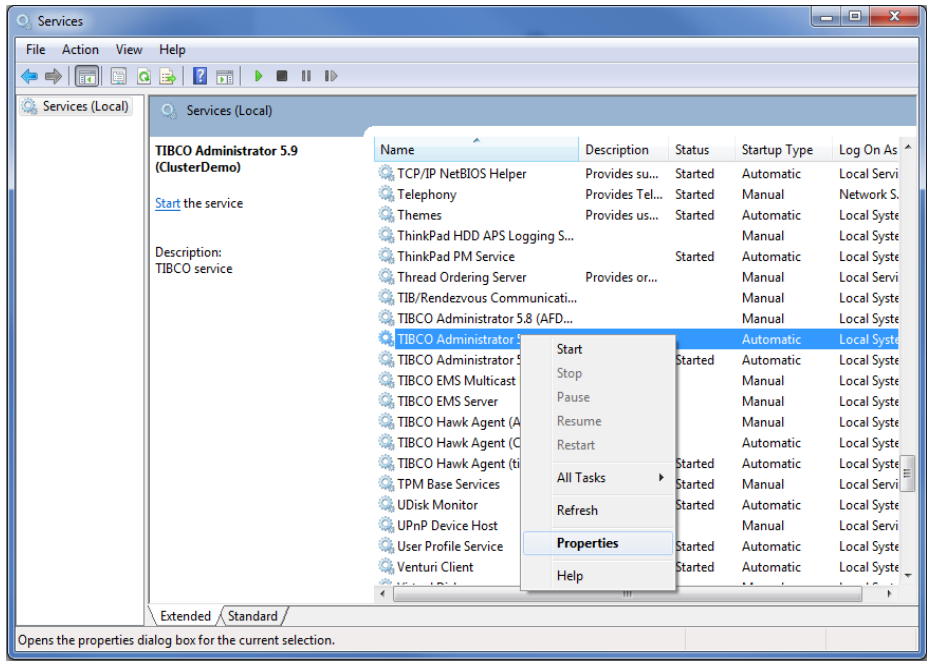
Figure 13 DomainUtility Confirm Window

10. Click Exit.

Register the Services with the Cluster

You must register the administration server and TIBCO Hawk Agent services with cluster software so it can be used to start and stop them. You'll need the service name when registering the service in the cluster. To get the service name, go to the Services panel, select the service name and click properties. For example:

Figure 14 Register the Services with the Cluster



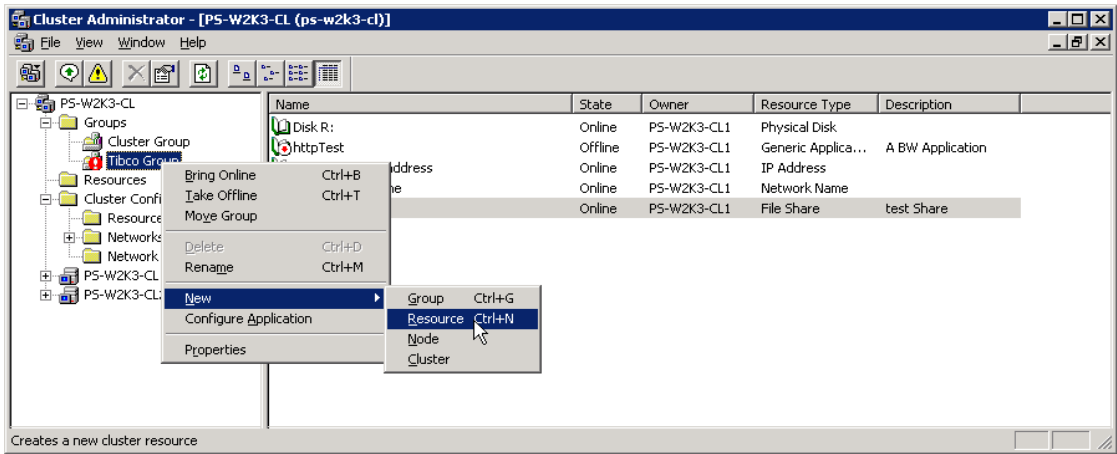
The name you'll need when registering the service with the cluster is listed in the Service Name field.

Add Service Name to the Cluster

You must complete this section twice. The first time to register the administration server service and a second time to register the TIBCO Hawk Agent service.

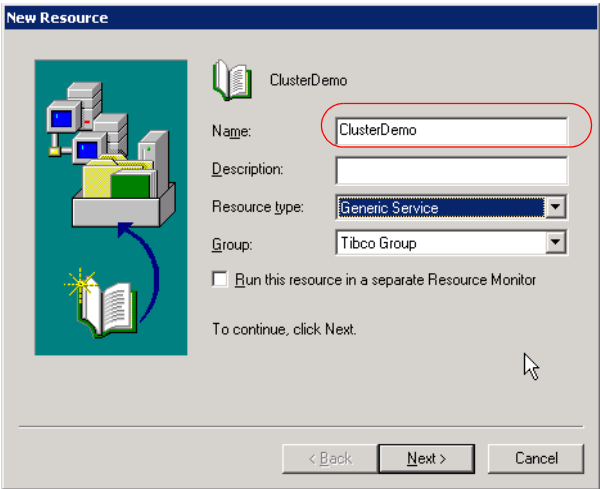
1. In the cluster administrator window, select **Tibco Group**, right-click and select **New > Resource**.

Figure 15 Cluster Administrator Window



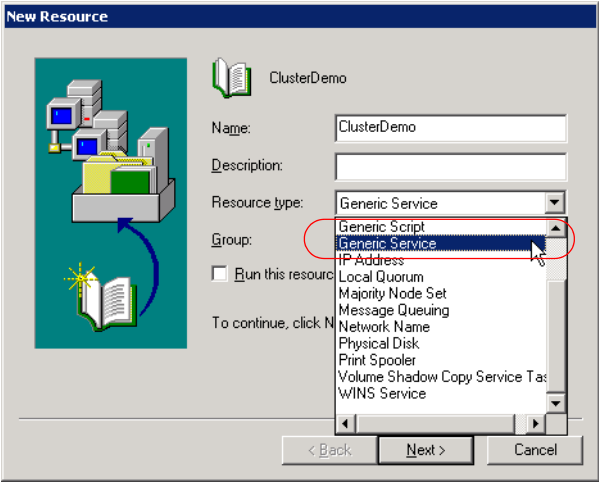
2. In the New Resource dialog, provide a name for the resource.

Figure 16 New Resource Dialog



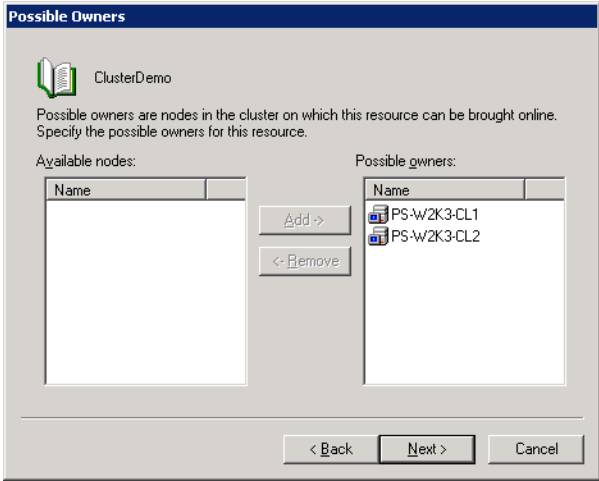
3. In the Resource Type field, select Generic Service.

Figure 17 Select Source Type



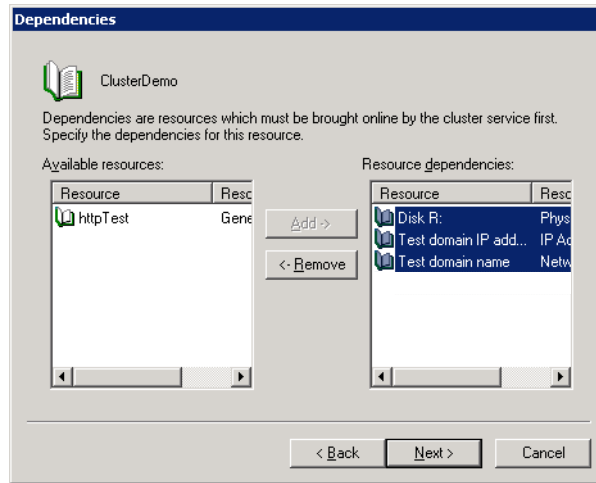
4. Click Next. In the Possible Owners dialog, the nodes that can own this resource should be listed in the Possible Owners column.

Figure 18 Possible Owners Dialog



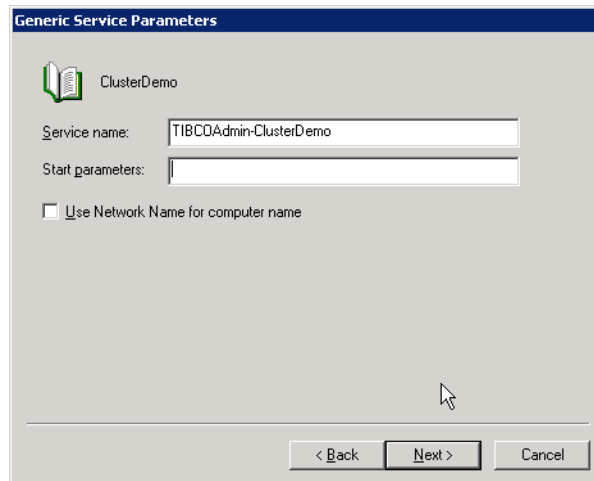
5. Click Next. In the Dependencies dialog, select the resources needed by the service.

Figure 19 Dependencies Dialog



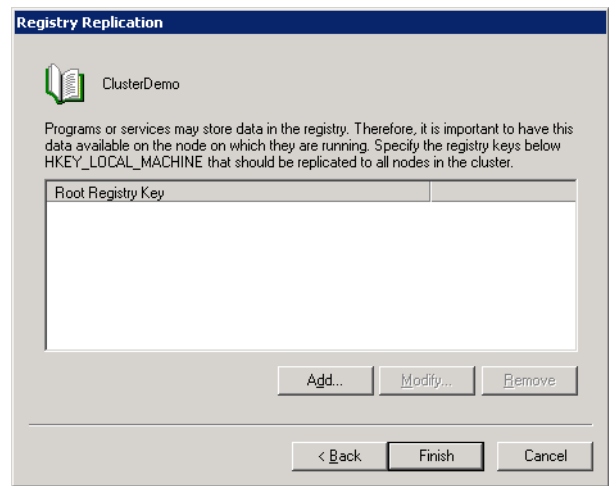
6. Click Next. Provide the service name you identified earlier.

Figure 20 Generic Service Parameters



7. Click Next. The Registry Replication dialog appears. No registry keys need be specified. Click Finish. A dialog appears indicating the resource was created.

Figure 21 Registry Replication



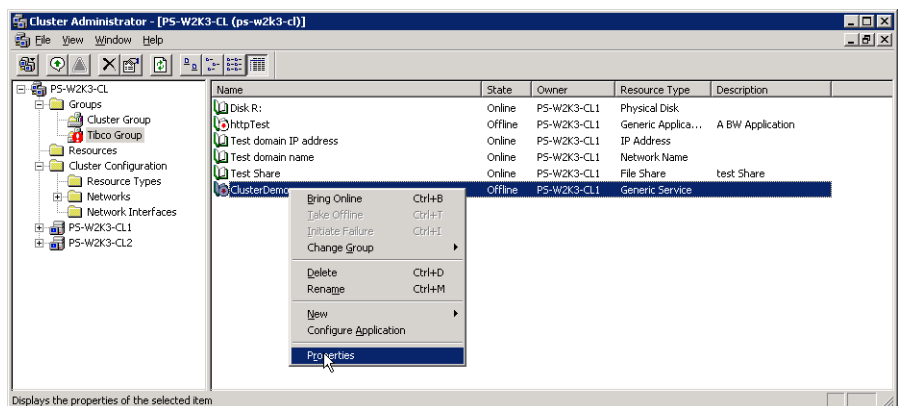
Change a Service Property

This section must be completed for the administration server service and the TIBCO Hawk Agent service.

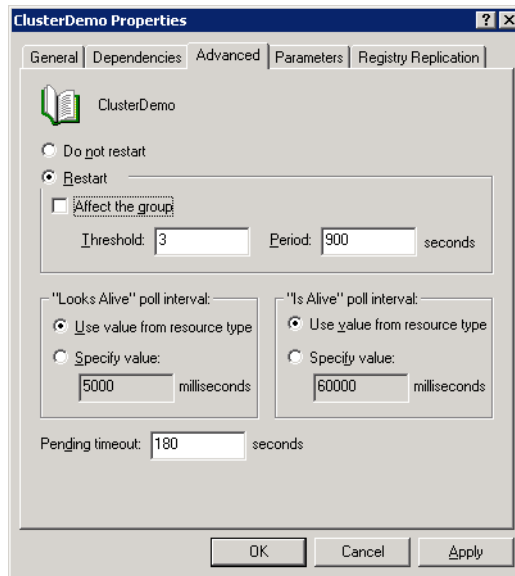
After registering the service, you must change a property set by default for the service.

- 1. In the cluster administrator window, right-click the service and select Properties.

Figure 22 Cluster Administrator Window



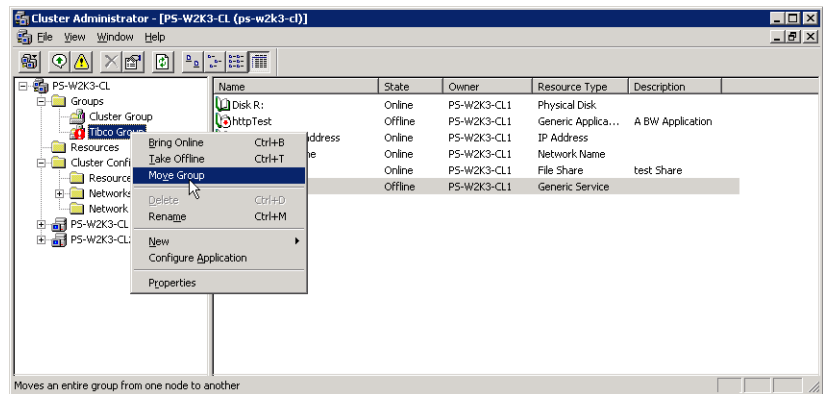
- 2. Click the Advanced tab and clear the Affect entire group check box. Click OK.

Figure 23 Advanced Tab

Configure the Second Node

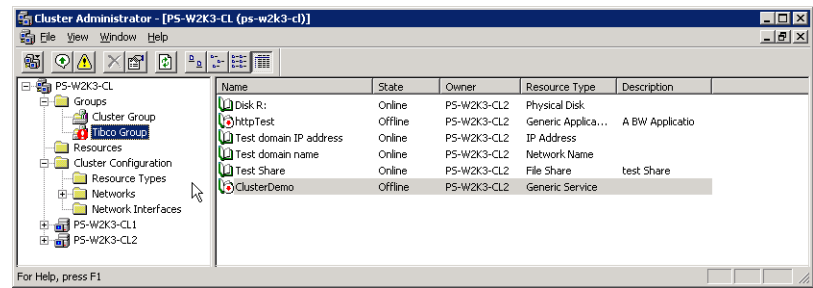
After installing TIBCO Runtime Agent, TIBCO Administrator, creating the initial domain on the first node, and registering the services in the cluster software, you must make the second node active. Right-click the group and select Move.

Figure 24 Right-click the Group and Select Move



When the move operation completes, the second node becomes active and you are ready to install TIBCO software on the node.

Figure 25 Active Second Move



Install TIBCO Software

After the second node is active, you must install TIBCO software on a local drive on the node (see [Installing TIBCO Software Onto a Shared Drive on page 46](#) for detailed instructions). Complete the steps that were explained in the following sections:

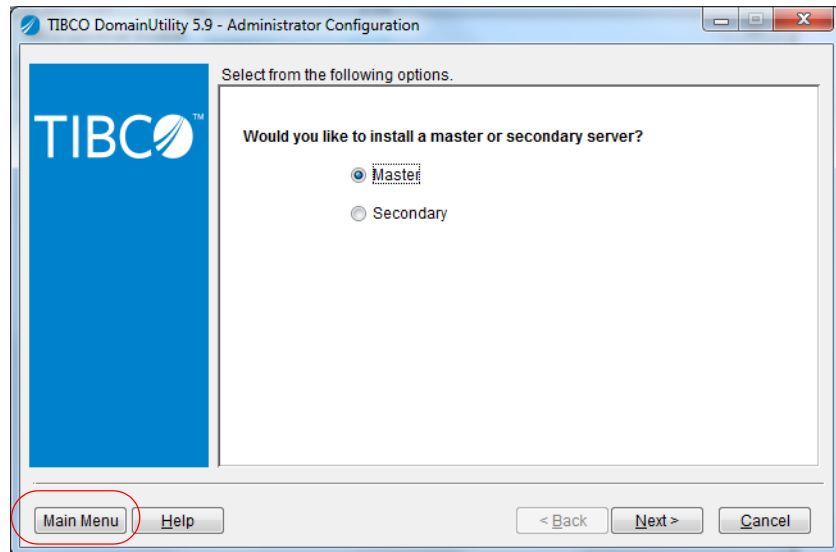
- [Installing TIBCO Runtime Agent on the Active Node on page 18](#)
- [Installing TIBCO Administrator on the Active Node on page 19](#)

Join the Second Node to the Domain

After installing TIBCO software, you must join the machine to the domain. The following steps explain how to do this.

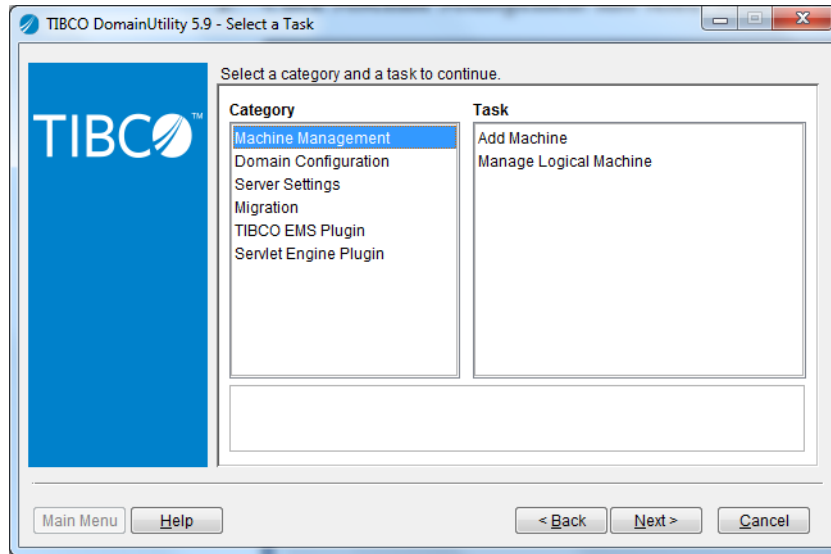
1. After installing TIBCO Administrator on the second node, TIBCO Domain Utility will be automatically launched to create the initial domain. Because the initial domain has already been created, click the Main Menu button, as shown in the next diagram.

Figure 26 Select Server




2. Click Machine Management and select Manage Logical Machine. Click Next.

Figure 27 Select a Task



3. Click Add. In the Configuration dialog, provide the domain name, logical machine name and the domain home location. The domain name and domain home was defined when the initial domain was configured for the first active node. The logical machine value is the cluster group network name. If a network name is not defined, provide the IP address. Click OK. For example, the next diagram shows the values defined for the installation example.

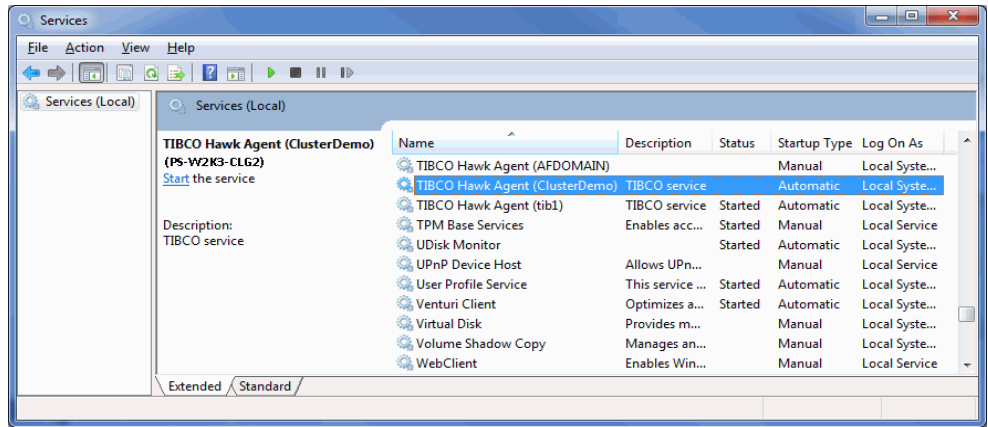
Figure 28 Configuration Dialog



4. Click Next, then click Exit.

After Domain Utility completes, you can verify that the service has been registered to the machine by navigating to the Services console. For example, the following entry displays in Services. Note that the service should not be started manually. Instead it will be registered with the cluster software and it will manage starting and stopping the service.


Figure 29 Services Console



Bring the Service Online

After the service is registered to the machine, you must use the cluster software to bring the service online. From the cluster administrator, right-click on the ClusterDemo generic resource and select Bring Online. The State column changes to online, indicating success.

Figure 30 Bring Online



Adding a TIBCO Application as a Service to the Cluster

This section shows how to register a TIBCO BusinessWorks process as a generic service to a cluster. The same steps apply if registering another TIBCO application as a service to a cluster.

You must install TIBCO BusinessWorks on the node before you can add a BusinessWorks process to the cluster. After installing the software, complete the following steps to add the process to the cluster. After the process has been added, you must use the cluster software to start and stop the process.

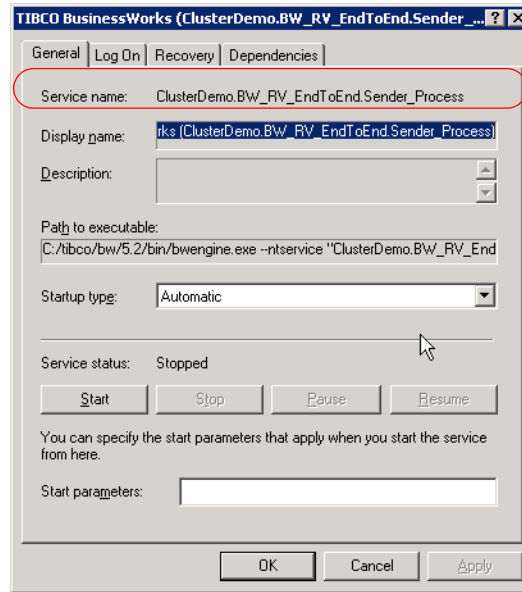
Add the Process as a Service

Before adding a process as a generic service to the cluster software, you must deploy the application that contains the process using the TIBCO Administrator GUI. If you are adding the process as a service to the cluster software, the process must be deployed as an NT Service.

Complete the following steps after the deployment is successful:

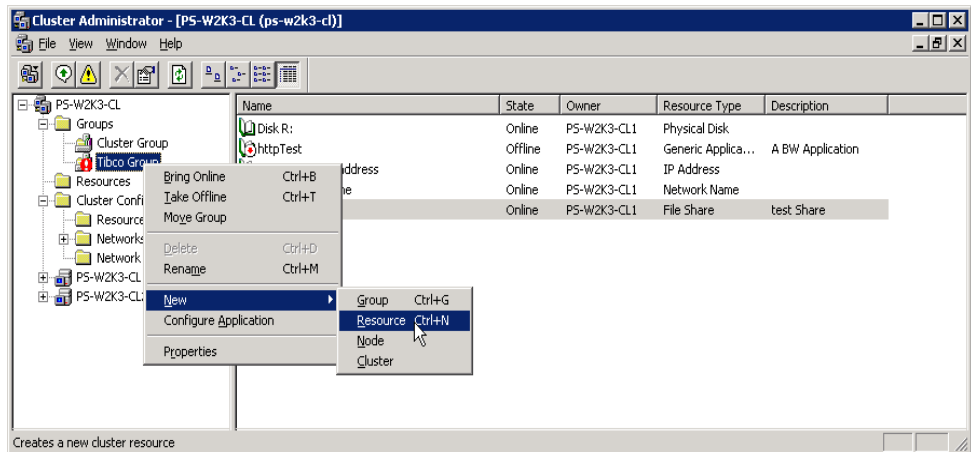
1. You'll need the service name when registering the BusinessWorks service in the cluster. To get the service name, go to the Services panel, click the service name and click properties. The name you'll need when registering the service with the cluster is listed in the Service Name field. For example:

Figure 31 Service Name Field



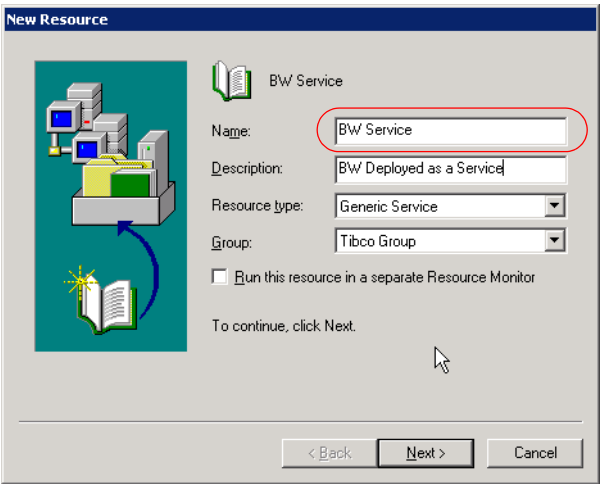
2. In the cluster administrator window, select **Tibco Group**, right-click and select **New > Resource**.

Figure 32 Cluster Administrator Window



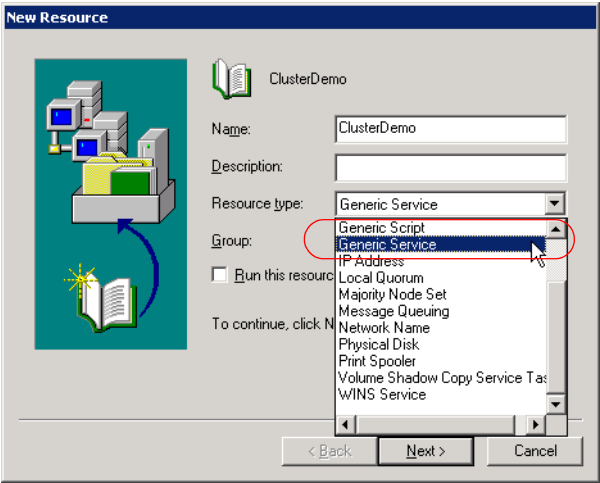
3. In the New Resource dialog, provide a name for the resource.

Figure 33 Provide Name for Resource



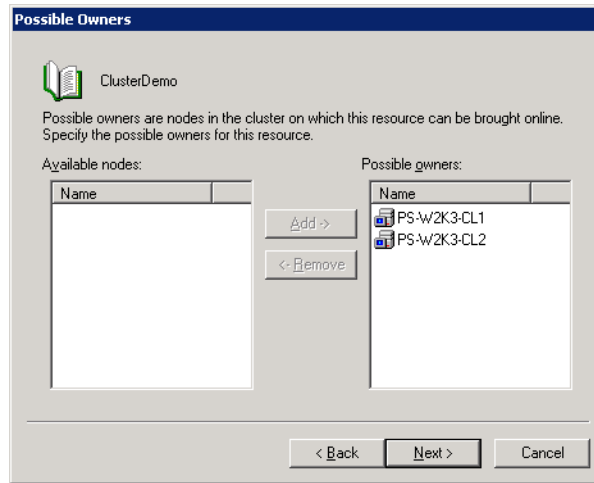
4. In the Resource type field, select Generic Service.

Figure 34 Resource Type Field



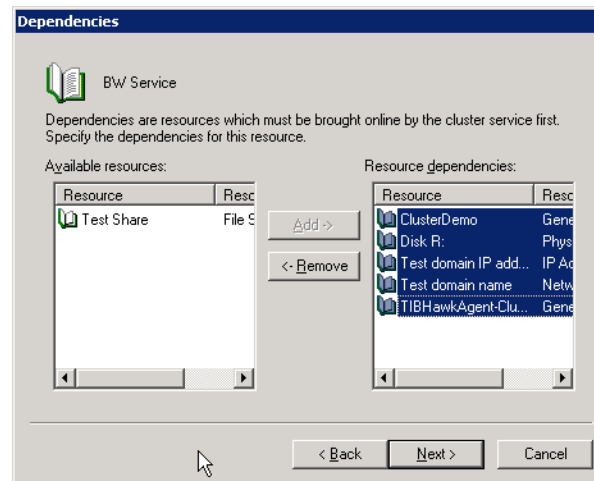
5. Click Next. In the Possible Owners dialog, the nodes that can own this resource should be listed in the Possible Owners column.

Figure 35 Possible Owners Window



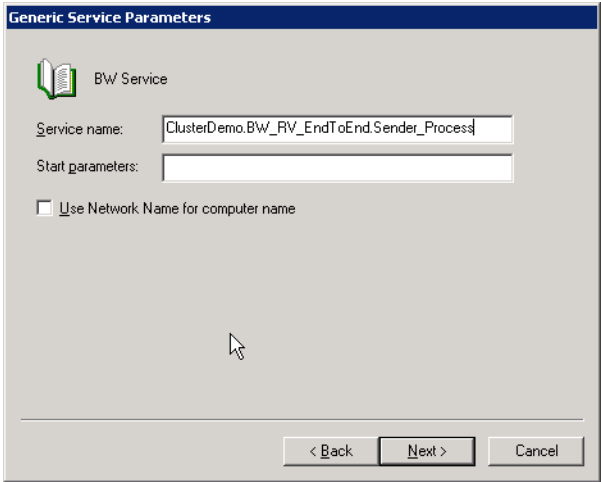
- Click Next. In the Dependencies dialog, select the resources needed by the service. This includes all the resources shown in the next diagram.

Figure 36 Dependencies Dialog



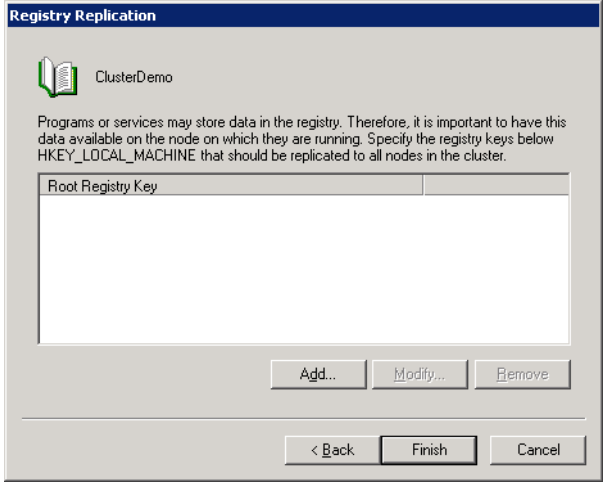
- Click Next. Provide the service name you identified in [step 1](#).

Figure 37 Enter Service Name



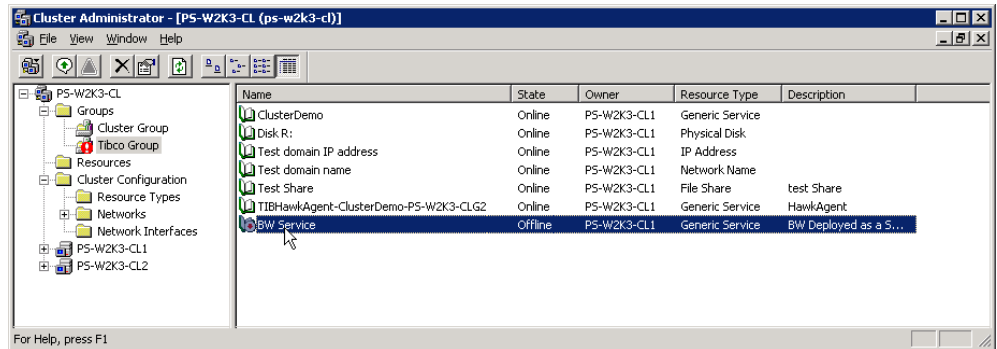
8. Click Next. The Registry Replication dialog appears. No registry keys need be specified. Click Finish.

Figure 38 Registry Replication Dialog



A dialog similar to the following appears indicating the resource was registered correctly.

Figure 39 Correct Registration of Resource

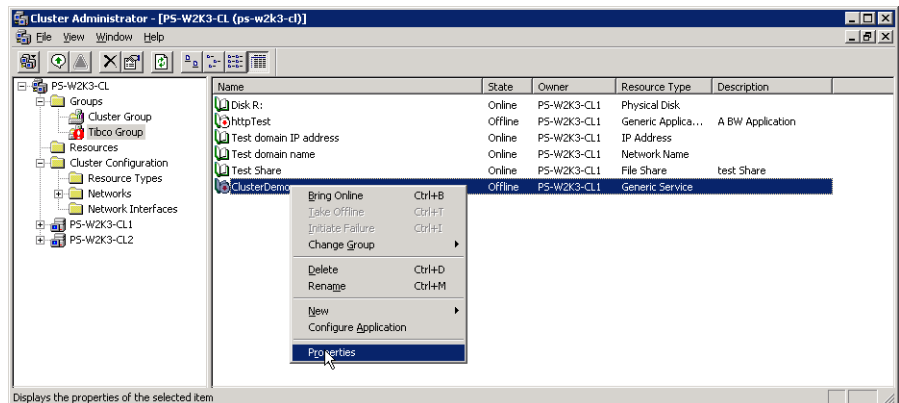


Change a Service Property

After registering the service, you must change a property set by default for the service.

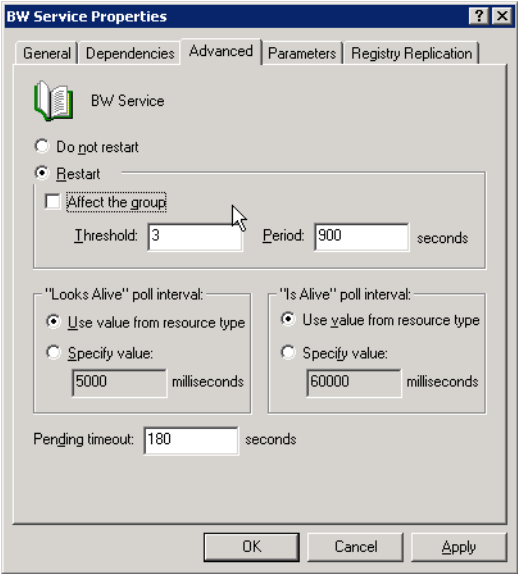
1. In the cluster administrator window, right-click the service and select Properties.

Figure 40 Cluster Administrator



2. Click the Advanced tab and clear the Affect entire group check box. Click OK.


Figure 41 BW Service Properties Window



Bring the Service Online

After the service is registered to the machine, you must use the cluster software to bring the service online. From the cluster administrator, right-click on the BW Service generic resource and select Bring Online. The State column changes to online, indicating success.

Figure 42 Bringing Service Online



Adding a TIBCO Application as an Application to a Cluster

This section shows how to register a TIBCO BusinessWorks process as a generic application to a cluster. The same steps apply if registering another TIBCO application as an application to a cluster.

You must install TIBCO BusinessWorks on the node before you can add a BusinessWorks process to the cluster. After installing the software, complete the following steps to add the process to the cluster. After the process has been added, you must use the cluster software to start and stop the process.

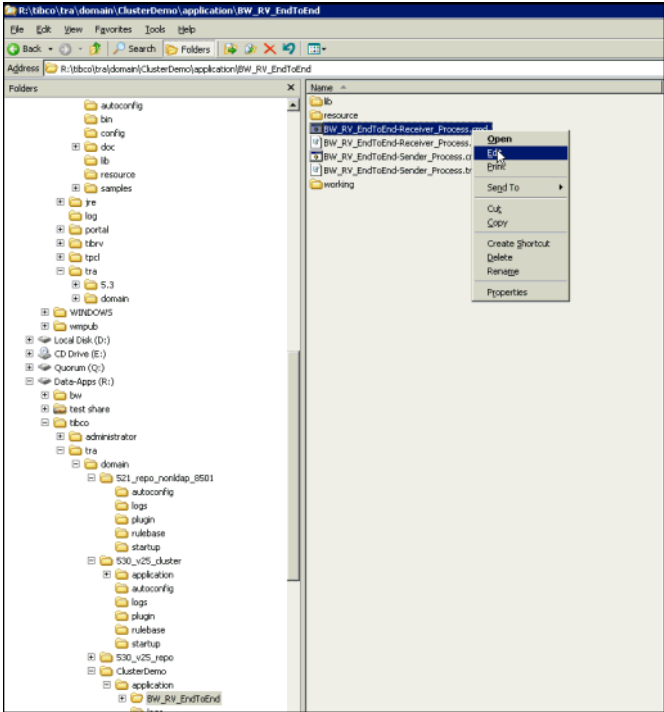
Add the Process as an Application

Before adding a process as a generic application to the cluster software, you must deploy the application that contains the process using the TIBCO Administrator GUI. If you are adding the process as an application to the cluster software, the process must be deployed as an application.

Complete the following steps after the deployment is successful:

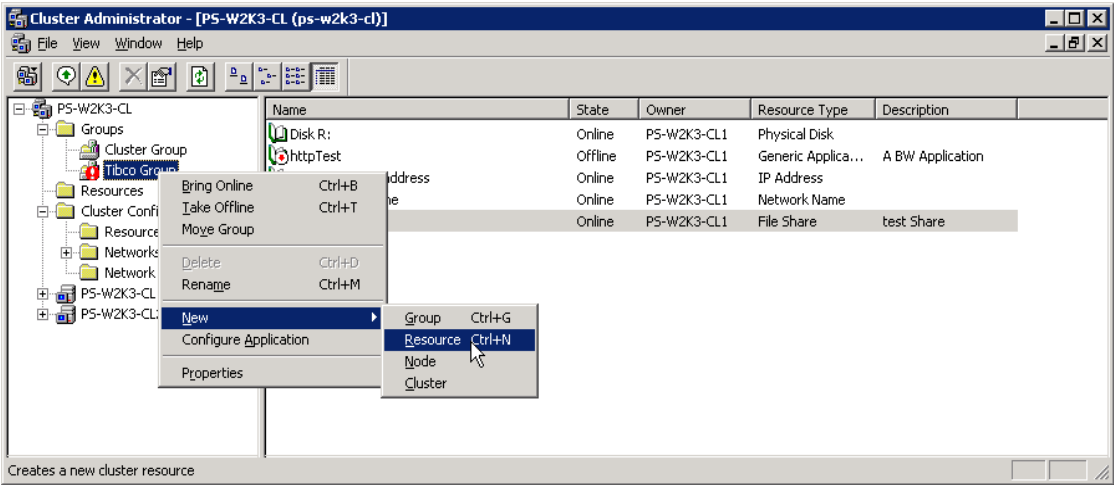
1. You'll need the application's command line when registering the BusinessWorks application to the cluster. To get the command line, navigate to the application directory on the cluster shared drive and open (by choosing Edit) the application's .cmd file. In our example, the application's .cmd file is located in the
`R:\tibco\tra\domain\ClusterDemo\application\BW_RV_EndtoEnd`
folder. Copy the entire command line. For example, see the following figure.

Figure 43 Application's .cmd File Location



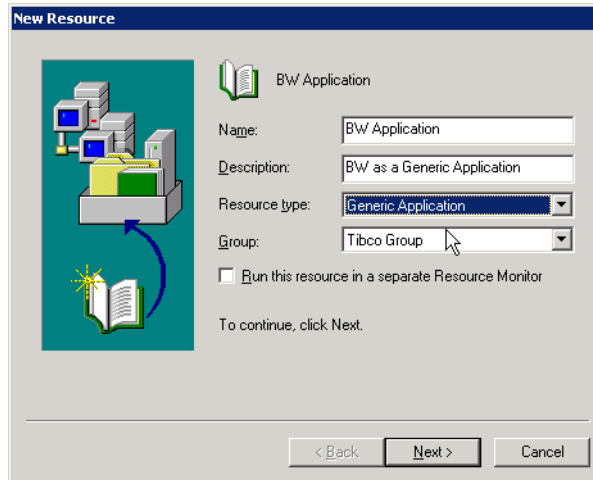
- 2. In the cluster administrator window, select **Tibco Group**, right-click and select **New > Resource**.

Figure 44 Cluster Administrator Window



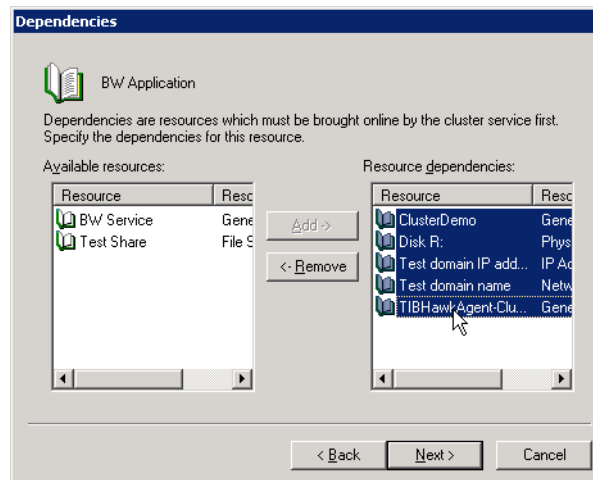
3. In the New Resource dialog, provide a name for the resource. In the Resource type field, select Generic Application.

Figure 45 Select Resource Type as Generic Application



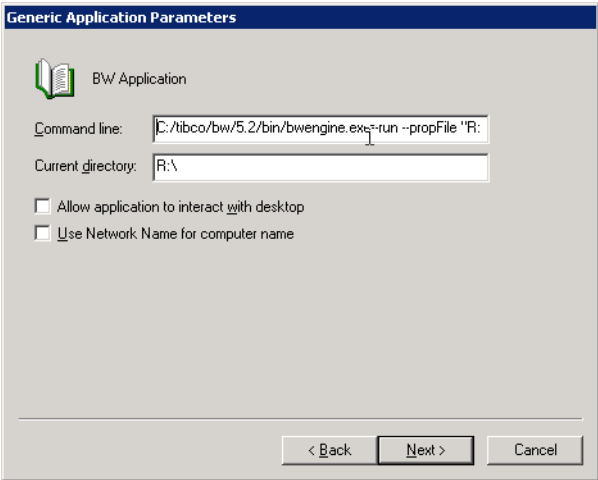
4. Click Next. In the Possible Owners dialog, the nodes that can own this resource should be listed in the Possible Owners column.
5. Click Next. In the Dependencies dialog, select the resources needed by the service. This includes all the resources shown in the next diagram.

Figure 46 Dependencies Dialog



- 6. Click Next. Paste the application’s command line that you copied in [step 1](#) in the Command Line field.

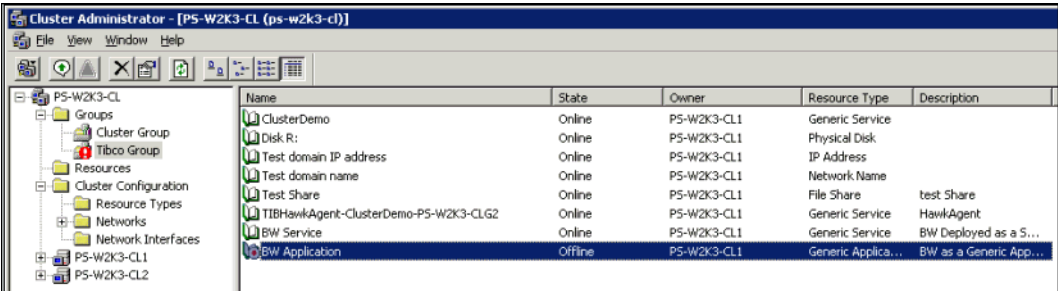
Figure 47 Generic Application Parameter Window



- 7. Click Next. The Registry Replication dialog appears. No registry keys need be specified. Click Finish.

A dialog similar to the following appears indicating the resource was registered correctly.

Figure 48 Correct ResourceRegistration



Change an Application Property

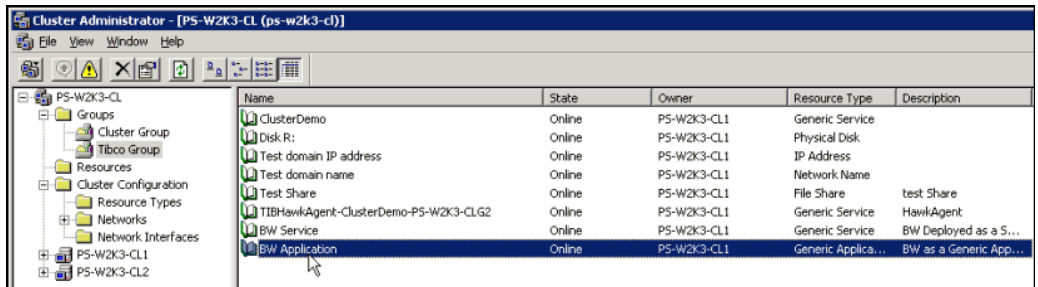
After registering the application, you must change a property set by default for the application.

1. In the cluster administrator window, right-click the application and select Properties.
2. Click the Advanced tab and clear the Affect entire group check box. Click OK.

Bring the Application Online

After the service is registered to the machine, you must use the cluster software to bring the service online. From the cluster administrator, right-click on the BW Application generic resource and select Bring Online. The State column changes to online, indicating success.

Figure 49 To Bring the Application Online



Installing TIBCO Software Onto a Shared Drive

You can install your TIBCO software on a shared drive, although it is not recommended. You only need to install the software on the first node, and copy-and-rename the installation registry file to the subsequent nodes.

To do so, follow the instructions below:

1. Map a network drive on the first node. Make sure that the same drive letter can be used on all the subsequent nodes.
2. Install all your TIBCO software on the first node.
3. For every subsequent node, map a network drive with the same drive letter assignment and network path reference as the one you mapped for the first node. For example, if you used X: as the drive letter and `//sample-fileserver/tibco-cluster/` as the network path for the drive, then you must do the same on all subsequent nodes.
4. Copy the TIBCO installation registry file (`vpd.properties.tibco.hostname`) from the first node to the `%SystemRoot%` directory in all the subsequent nodes and rename the hostname accordingly.

For Windows machines, the `%SystemRoot%` directory is `C:\Windows`. See *TIBCO Runtime Agent Installation* for more information on the TIBCO installation registry files.



If you wish to install additional TIBCO software in the cluster (TIBCO BusinessWorks, for example), you must install it on the first node and copy-and-rename the installation registry file to the other nodes again.

Example

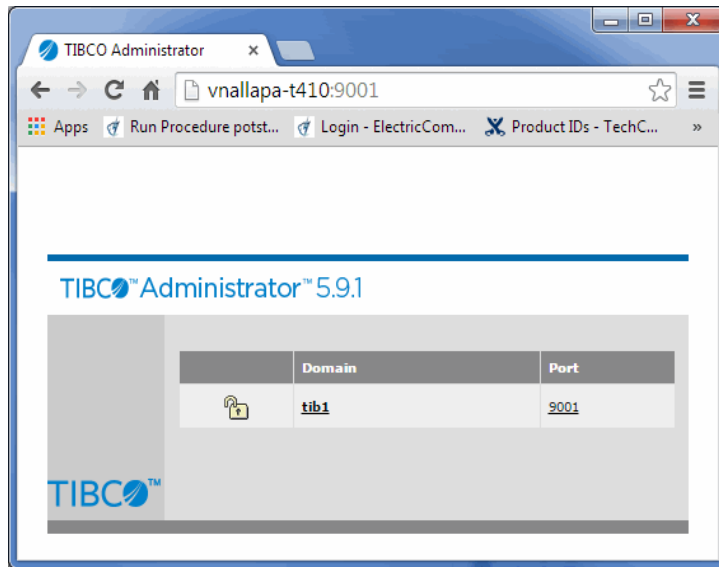
In the example cluster given in this chapter, do the following to install TIBCO Runtime Agent and TIBCO Administrator on a shared drive:

1. On PS-W2K3-CL1, map a network drive to `//PS-W2K3-CLSHARE/D` and assign it the letter X.
2. Install TIBCO Runtime Agent and TIBCO Administrator on PS-W2K3-CL1 in the `X:\tibco` directory. See *TIBCO Runtime Agent Installation* and *TIBCO Administrator Installation* for detailed instructions.
3. On PS-W2K3-CL1, copy the file `C:\Windows\vpd.properties.tibco.ps-w2k3-cl1`.
4. Paste the copied file into `C:\Windows` on PS-W2K3-CL2 and rename it to `vpd.properties.tibco.ps-w2k3-cl2`.

Verify the Installation

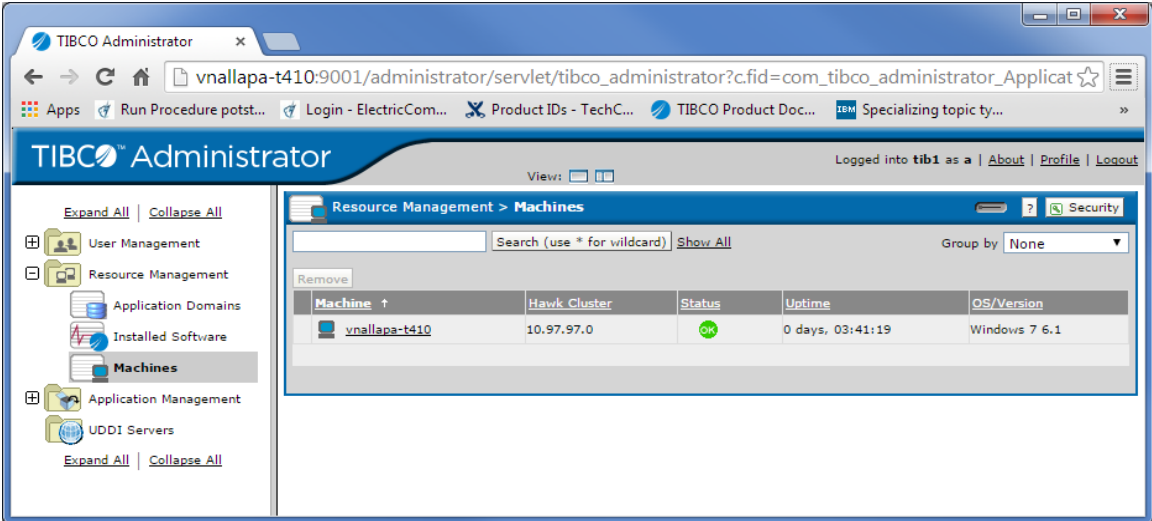
You can verify the cluster installation by starting TIBCO Administrator. You must use the cluster network name to display the TIBCO Administrator GUI login screen. In the example, the login address is `http://ps-w2k3-clg2:8090`. This is shown in the next diagram.

Figure 50 TIBCO Administrator GUI Login Screen



After logging in, under Resource Management > Machines the following dialog displays. Note that the machine icon indicates the machine is a virtual machine.

Figure 51 Virtual Machine



Installed Files

This section explains the directory structure and files installed on each node in the cluster. The files are used to manage a domain that is part of a cluster environment.

Directory Structure

The following directory structure is added to the existing TIBCO Administrator domain directory structure. The `tibcoadmin_domain.tra`, `tibcoadmin_domain.exe` and `server_domain.xml` files are created in these directories. This allows all TIBCO Administrator domain-specific information to be isolated on a shared drive for clustering purposes.

```
TIBCO_ADMIN_DOMAIN_HOME\domain-name\bin
TIBCO_ADMIN_DOMAIN_HOME\domain-name\tomcat\conf
TIBCO_ADMIN_DOMAIN_HOME\domain-name\tomcat\logs
```

Files

The *machine* value is added to the `TIBCO_TRA_DOMAIN_HOME` entry only if a domain was created with a virtual machine name or a virtual machine was added to a domain. This is necessary to differentiate among multiple entries, since a physical machine can be identified through multiple virtual machine names and each virtual machine can added to the same domain.

The `tibcoadmin_domain.tra` file created includes the following properties:

```
java.property.TIBCO_ADMIN_DOMAIN_HOME=location
java.property.TIBCO_TRA_DOMAIN_HOME=location
```

Domain Home Tracking

TIBCO Domain Utility creates the `DomainHomes.properties` file to track all the TIBCO Runtime Agent and TIBCO Administrator domain homes on the machine. If the TIBCO Runtime Agent home is at the location, `C:\tibco\tra\<version>`, the file is created at the directory level: `C:\tibco\tra\domain`.

Chapter 3

Deploying and Starting Applications Under Cluster Control

This chapter explains how applications under cluster control should be deployed, started and stopped.

Topics

- [Redeploying Applications, page 52](#)
- [Forcing Redeployment of an NT Service, page 53](#)
- [Starting and Stopping Applications Managed by Cluster Software, page 54](#)
- [Removing Orphaned Services, page 55](#)

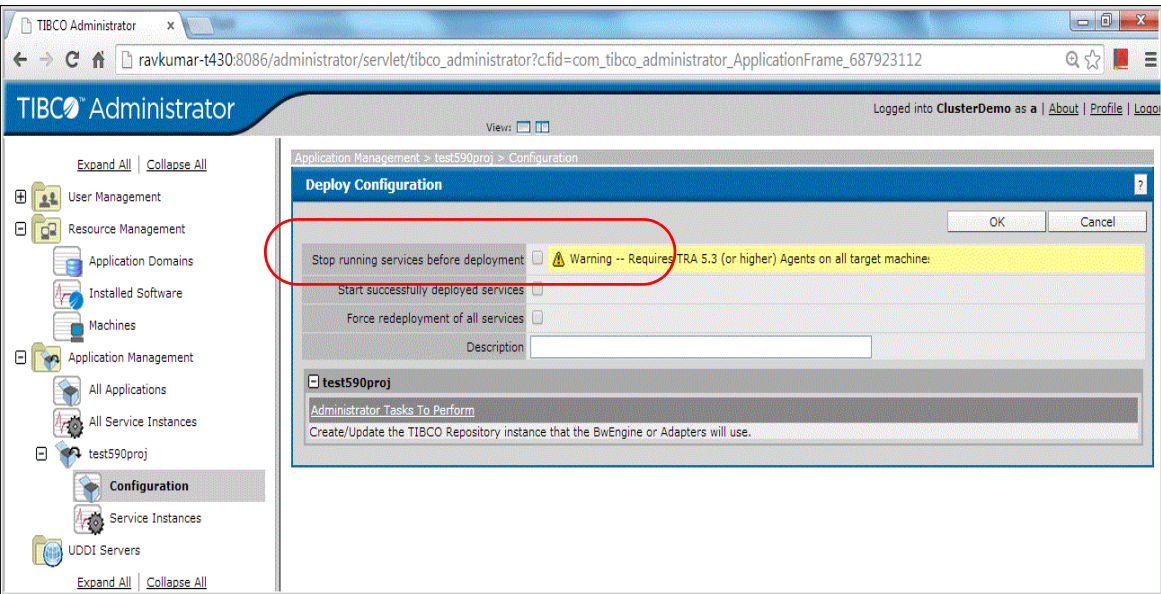
Redeploying Applications

If a service is registered to the cluster software as auto-restarted if failed, when the service is redeployed using the TIBCO Administrator GUI, the Stop running service before deployment option should be cleared (as shown in the next diagram).

By default, redeployment will stop the running instance before it is redeployed. If the option is cleared, the running instance will not be stopped during redeployment, which prevents the cluster software from mistakenly detecting that the instance has failed and prevents the cluster software from restarting the instance during redeployment.

After the application is redeployed, you must use the cluster administrator software to start and stop the application.

Figure 52 Using Cluster Administrator Software to Start and Stop the Application



Forcing Redeployment of an NT Service

Under the Microsoft Cluster Server, an NT service can be registered as generic service. When you deploy a service as an NT service from the TIBCO Administrator GUI, an NT service entry is created on the active cluster node that is hosting the cluster group at that moment. However, the NT service entry is not created on inactive cluster nodes.

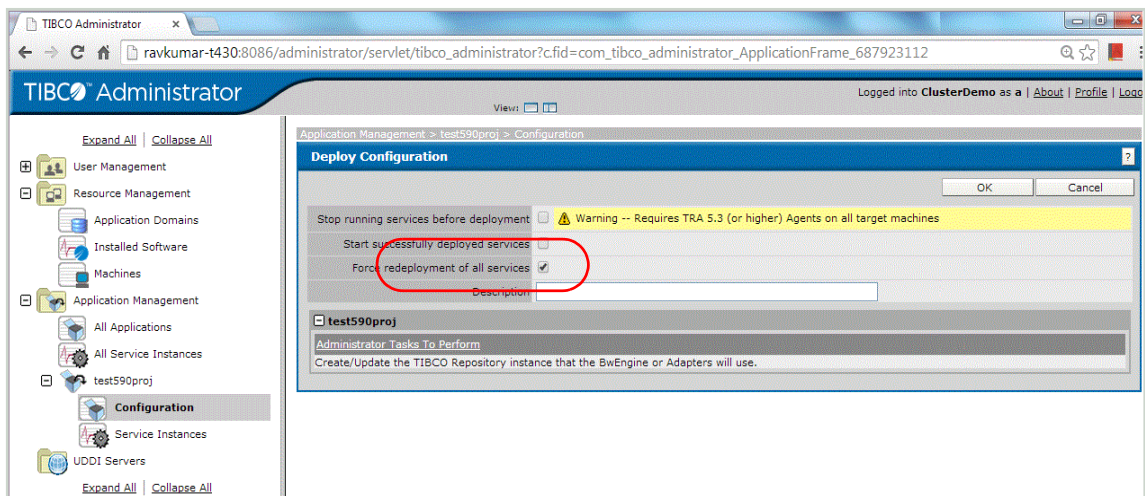
After an inactive node becomes active, applications running as NT services must be deployed on the newly activated node. Applications running as NT services must also be deployed when an inactive node becomes active if parameters for the application are changed.

To duplicate the NT service entry on an inactive cluster node, you must move the cluster group from the currently active node to the inactive node (making it active), and then in the TIBCO Administrator GUI, select the Force redeployment of all services option to redeploy the already successfully deployed application.



If the application is not running as an NT Service, you need not deploy the application.

Figure 53 Force Redeployment



Starting and Stopping Applications Managed by Cluster Software

If an application is managed by cluster software, it is recommended to not start or stop it from the TIBCO Administrator GUI. This applies both to applications running as an NT Service and applications not running as a service. For example, if a TIBCO BusinessWorks process is managed using cluster software, the process should be started and stopped using the cluster software, not using the TIBCO Administrator GUI.

Removing Orphaned Services

If a service is undeployed, its NT service entry is only removed from the currently active cluster node. The other cluster nodes will have orphaned service entries. To remove the orphaned service entries, you must execute the following command on the inactive cluster node:

```
wrap --ntservice nt service name --uninstall
```


Index

A

- Active/Active [4](#)
- Active/Passive [4](#)
- Adding a TIBCO Application as a Service to the Cluster [34](#)
 - Add the Process as a Service [34](#)
 - Bring the Service Online [40](#)
 - Change a Service Property [39](#)
- Adding a TIBCO Application as an Application to a Cluster [41](#)
 - Add the Process as an Application [41](#)
 - Bring the Application Online [45](#)
 - Change an Application Property [44](#)

C

- cluster [2, 2](#)
- cluster environment [2](#)
- Cluster Support Limitations [7](#)
- Cluster Terminology [3](#)
 - Active/Active [4](#)
 - Active/Passive [4](#)
 - Dependency [4](#)
 - failback [4](#)
 - Failover [4](#)
 - Group [3](#)
 - Node [3](#)
 - Resource [3](#)
 - Rolling Upgrade [4](#)
 - Shared storage [4](#)
- cluster-aware applications [2](#)
- Cold standby [3](#)
- Configure the Second Node [30](#)

Bring the Service Online [33](#)

Install TIBCO Software [31](#)

Join the Second Node to the Domain [31](#)

Creating a Cluster Group [17](#)

customer support [xvi](#)

D

- Dependency [4](#)
- Deploying and Starting Applications Under Cluster Control [51](#)
- Domain Not in Cluster and Services in Cluster [11](#)

E

- Entire Domain in a Single Cluster Group [9](#)
- Entire Domain Split Into Multiple Cluster Groups [10](#)
- ENV_NAME [xiii](#)
- Example Installation Diagram [15](#)

F

- failback [4](#)
- Failover [4](#)
- Failover Cluster [5](#)
- Forcing Redeployment of an NT Service [53](#)

G

- Group [3](#)

H

High-availability Modes [3](#)

 Cold standby [3](#)

 Hot standby [3](#)

 Warm standby [3](#)

Hot standby [3](#)

I

Installed Files [49](#)

 Directory Structure [49](#)

 Domain Home Tracking [49](#)

 Files [49](#)

Installing the Software Into a Cluster [13](#)

Installing TIBCO Administrator on the Active

 Node [19](#)

 Create the Initial Domain [19](#)

Installing TIBCO Runtime Agent on the Active

 Node [18](#)

Installing TIBCO Software Onto a Shared Drive [46](#)

 Example [46](#)

L

luster-unaware applications [2](#)

N

Node [3](#)

P

place TIBCO applications into a cluster scenarios [8](#)

R

Redeploying Applications [52](#)

Register the Services with the Cluster [24](#)

 Change a Service Property [28](#)

 Add Service Name to the Cluster [24](#)

Removing Orphaned Services [55](#)

Resource [3](#)

Rolling Upgrade [4](#)

S

Shared storage [4](#)

Starting and Stopping Applications Managed by Cluster Software [54](#)

support, contacting [xvi](#)

T

technical support [xvi](#)

TIBCO_HOME [xiii](#)

V

Verify the Installation [47](#)

W

Warm standby [3](#)