

TIBCO® Data Virtualization

Active Cluster Guide

Version 8.0

Last Updated: November 7, 2018

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 and the TIBCO logo are either registered trademarks or trademarks of TIBCO Software Inc. in the United States and/or other countries

TIBCO, Two-Second Advantage, TIBCO Spotfire, TIBCO ActiveSpaces, TIBCO Spotfire Developer, TIBCO EMS, TIBCO Spotfire Automation Services, TIBCO Enterprise Runtime for R, TIBCO Spotfire Server, TIBCO Spotfire Web Player, TIBCO Spotfire Statistics Services, S-PLUS, and TIBCO Spotfire S+ are either registered trademarks or trademarks of TIBCO Software Inc. in the United States and/or 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.

Copyright © 2002-2018 TIBCO Software Inc. All rights reserved.

TIBCO Software Inc. Confidential Information

Contents

Preface	7
Product-Specific Documentation	7
How to Access TIBCO Documentation	8
How to Contact TIBCO Support	8
How to Join TIBCO Community	8
Introducing Active Cluster	9
Overview of Active Cluster	9
Advantages of Active Cluster	10
Active Cluster Architecture	11
Metadata Repository	12
Cache Data Sharing	13
Using a Load Balancer	13
About the Timekeeper	14
Triggers	15
Propagation of Metadata Changes	15
Error Recovery	15
Joining and Leaving an Active Cluster	15
Security Implementation	16
Installation	16
Planning the Deployment of Active Cluster	17
Considerations for Deploying Active Cluster	17
Deployment Environments	17
Determining the Number of Nodes in a Cluster	18
Performance Considerations	18
Monitor Considerations	18
Caching	18
Deployment Example	19
Typical Process for Deployment	19
Working with Active Cluster	21
About Working with Active Cluster	21
About Manager with Active Cluster	22
About the cluster_util Program	23
About Studio with Active Cluster	23
Privileges Required for Working with Active Cluster	23

Triggers and Active Cluster	24
prePropagation of Custom Adapters in an Active Cluster	25
Creating a New Active Cluster	25
Adding a TDV Server to an Active Cluster	27
Preparing to Join an Active Cluster	28
Joining an Active Cluster	29
Renaming an Active Cluster	32
Removing a TDV Server from an Active Cluster	34
Propagate a Domain Between Clusters	36
Delete an Active Cluster	36
Viewing the Status of an Active Cluster	36
Getting the Active Cluster and Member Names	39
Configuring an Active Cluster	40
Configuring Logging for Active Cluster	43
Changing the Port of a Cluster Node	44
Handling Active Cluster Errors	47
Overview of Error Handling	47
Setting a Timekeeper and Repairing Active Cluster Nodes	48
Regroup Using Manager	48
Regroup Using the cluster_util Script	49
Regroup Using the API	50
Troubleshooting	50
Issues with Setting Up a Cluster	51
A Cluster Node Is DISCONNECTED	51
A Cluster Node Goes Offline	51
A Cluster Node Is Busy and Does Not Respond	51
Nodes Have Been Partitioned into Subgroups	52
A Node Has Been Evicted Due to Metadata Conflicts	52
A Cluster Node Is Frequently DISCONNECTED	53
The Timekeeper Cluster Node Fails or Is Taken Offline	53
A Cluster Node Is Frequently BLOCKED	54
Cannot start CMS after Backing Up a Cluster Node	54
Other Best Practices for Active Cluster	55
Work with Large Amounts of Data	55
Work with the Command Line Interface	55
The cluster_util Command-Line Program	57
Using the cluster_util Program	57
The -create Subcommand	60

The -getConfig Subcommand61

The -join Subcommand61

The -remove Subcommand62

The -setClusterName Subcommand62

Active Cluster SYS_CLUSTER System Table63

Preface

Documentation for this and other TIBCO products is available on the TIBCO Documentation site. This site is updated more frequently than any documentation that might be included with the product. To ensure that you are accessing the latest available help topics, please visit:

- <https://docs.tibco.com>

Product-Specific Documentation

The following documents form the TIBCO® Data Virtualization(TDV) documentation set:

- *TIBCO TDV and Business Directory 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.
- TDV Installation and Upgrade Guide
- TDV Administration Guide
- TDV Reference Guide
- TDV User Guide
- TDV Security Features Guide
- Business Directory Guide
- TDV Application Programming Interface Guide
- TDV Tutorial Guide
- TDV Extensibility Guide
- TDV Getting Started Guide
- TDV Client Interfaces Guide
- TDV Adapter Guide
- TDV Discovery Guide
- TDV Active Cluster Guide
- TDV Monitor Guide
- TDV Northbay Example

How to Access TIBCO Documentation

Documentation for TIBCO products is available on the TIBCO Product Documentation website mainly in the HTML and PDF formats.

The TIBCO Product Documentation website is updated frequently and is more current than any other documentation included with the product. To access the latest documentation, visit <https://docs.tibco.com>.

Documentation for TIBCO Data Virtualization is available on <https://docs.tibco.com/products/tibco-data-virtualization-server>.

How to Contact TIBCO Support

You can contact TIBCO Support in the following ways:

- For an overview of TIBCO Support, visit <https://www.tibco.com/services/support>.
- For accessing the Support Knowledge Base and getting personalized content about products you are interested in, visit the TIBCO Support portal at <https://support.tibco.com>.
- 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 <https://support.tibco.com>. 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](https://community.tibco.com). For a free registration, go to <https://community.tibco.com>.

Introducing Active Cluster

This topic describes the features and architecture of Active Cluster.

- [Overview of Active Cluster, page 9](#)
- [Advantages of Active Cluster, page 10](#)
- [Active Cluster Architecture, page 11](#)
- [Security Implementation, page 16](#)
- [Installation, page 16](#)

Overview of Active Cluster

Active Cluster is a TIBCO® Data Virtualization (TDV) product option that allows you to deploy TDV in active/active clustered environments.

An Active Cluster comprises a group of TDV instances connected to each other over a system network. Active Cluster guarantees that all nodes are identical, with their metadata activity and configuration information automatically synchronized in the background.



Active Cluster

To maximize scalability and availability, deploy Active Cluster with a network load balancer.

Each cluster node is a complete TDV Server with its own repository. When a TDV Server joins an Active Cluster, it is automatically synchronized with the others.

Advantages of Active Cluster

Active Cluster helps IT organizations provide:

- High availability
- Scalability
- Reduced RPO (Recovery Point Objective) and RTO (Recovery Time Objective)

High Availability

High availability in Active Cluster is achieved by ensuring that all TDV cluster nodes are redundant: if any node fails or is removed from the cluster, the remaining nodes continue servicing requests without administrative intervention.

Active Cluster provides:

- With properly planned clusters, active/active clustering achieves 99.999% availability. Because all nodes are identical, traffic intended for a failed node can be automatically passed to another node.
- When a failed node comes back on line, it automatically rejoins the cluster and is resynchronized.
- All nodes are notified of node failures.

You can view the status of the Active Cluster and its members in the Manager Web interface. If a cluster node is out of sync with the other nodes, Active Cluster offers options for resolving the issue.

Scalability

Active Cluster helps your organization scale the TDV environment on demand, based on changing requirements:

- You can dynamically add cluster nodes to address increased load.
- Cluster nodes are automatically synchronized to maintain identical metadata.
- The cluster, available through a load balancer, appears to clients to be a single server.

Reduced RPO and RTO

Active Cluster helps to ensure that your RPO and RTO objectives are met. With Active Cluster:

- Repository can be enabled for HA/DR (High Availability/Disaster Recovery) through native or third-party capabilities.
- Migration and recovery of metadata occurs automatically, minimizing impact on TDV.

Active Cluster Architecture

An Active Cluster is initially set up on a computer running TDV. Additional computers (nodes) running TDV can then join the Active Cluster. Each cluster node is identical to a standalone TDV server in terms of system configuration and repository support.

All cluster nodes must be licensed and running the same version and patch level of Active Cluster. See the *TDV Installation and Upgrade Guide* for more information.

Active Cluster uses a share-nothing architecture, so you can mix UNIX and Windows servers in a cluster. However, we recommend that all nodes be of comparable operational characteristics: memory, number of CPUs, and CPU speed.

All servers are peers. Caching, statistics, and triggers are executed in a round-robin sequence.

Active Cluster characteristics:

- [Metadata Repository, page 12](#)
- [Cache Data Sharing, page 13](#)
- [Using a Load Balancer, page 13](#)
- [About the Timekeeper, page 14](#)
- [Triggers, page 15](#)
- [Propagation of Metadata Changes, page 15](#)
- [Error Recovery, page 15](#)
- [Joining and Leaving an Active Cluster, page 15](#)
- [Security Implementation, page 16](#)

Metadata Repository

Each TDV server has its own metadata repository, which can reside on the server or in a database instance. All metadata is replicated except for information specific to each server.

Clustered artifacts include:

- Global configuration settings
- Studio resources, including: containers, data sources, definition sets (SQL and XML), published resources, Discovery models, procedures (SQL Script, Packaged Query, Xquery, XSLT), tables (relational tables, views, flat files), trees (XML files), and triggers.
- Domain, group, and user information
- Primary CJP JAR files
- Custom adapter files
- Data source statistics
- Designer bundles

TDV metadata that is not replicated includes:

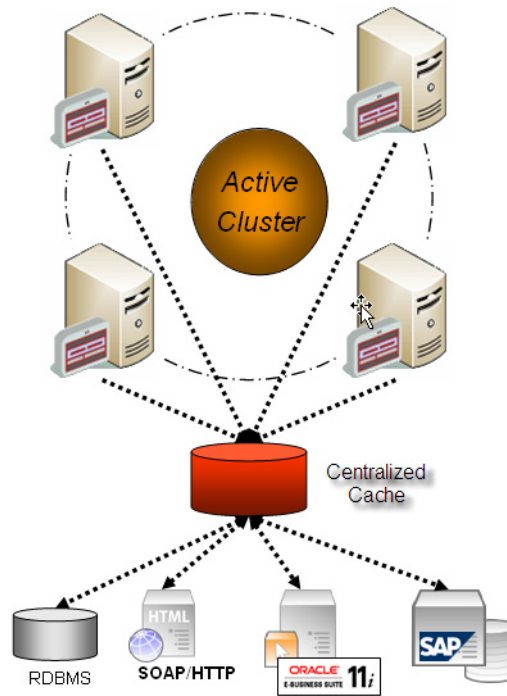
- Local configuration settings
- TDV instance-bound information:
 - Runtime information
 - File cache data
- Data requiring manual intervention:
 - Data source files (for example, XML and CSV)
 - Resources referenced by primary CJP JAR files
 - Customized ldap.properties files

All non-replicating configuration parameters are marked in the Administration > Configuration window with **This is not replicated in a cluster** in bold in the Description area.

Customized ldap.properties files and data source capabilities files are not automatically replicated or synchronized; these must be replicated manually. See *TDV Installation and Upgrade Guide* for more information.

Cache Data Sharing

Active Cluster uses a single, centralized cache to reduce the impact on back-end database systems. The first node to access a data source caches the data. Subsequent calls from other nodes access the cached data.



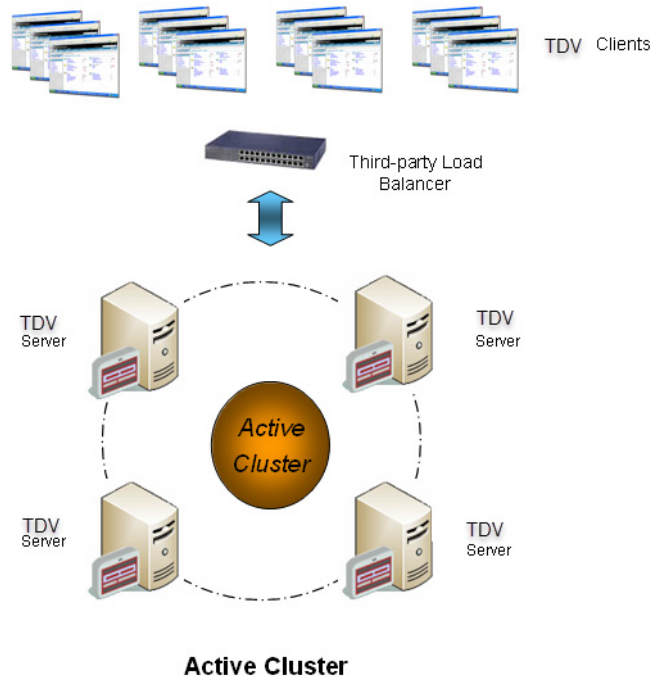
All servers in the cluster know the state and location of all database caches, regardless of which server does the refresh, because this information is stored in the database. TDV Manager provides a cluster status table that displays which server performed the refresh.

A file-system cache cannot be shared by cluster servers because status is stored on the local file system. File-system caches are replicated on every node in the cluster.

Using a Load Balancer

Clients typically connect to a third-party load balancer that forwards each connection request to the next available node. However, Active Cluster does not require a load balancer and it has no programmatic awareness of the load balancer. All configuration with a load balancer is performed outside of Active Cluster.

A typical scenario is a four-node Active Cluster with a third-party load balancer.



About the Timekeeper

One TDV cluster node serves as a timekeeper. The timekeeper generates a sequence number for each requested metadata update to the cluster. Changes are then made in FIFO order. The timekeeper also makes sure that triggers operate as configured. Finally, the timekeeper regulates how nodes join or leave the cluster. The timekeeper is otherwise identical to all other TDV server nodes.

Active Cluster automatically determines which TDV Server is the timekeeper. When a cluster is created, the first member is made the timekeeper. If the timekeeper machine goes down, another TDV Server node automatically assumes the timekeeper duties.

Because the timekeeper node performs approximately the same workload as other cluster nodes, it is not necessary to maintain a special server as the timekeeper node.

Triggers

Triggers are controlled by the timekeeper. Active Cluster automatically handles timer-event triggers. By default, a timer-event trigger executes exactly once per cluster, and other nodes are notified.

Non-timer event triggers are executed on each node in the cluster.

Propagation of Metadata Changes

Once applied successfully on the local node, each metadata change is propagated to all other nodes in parallel before the next change is propagated. Therefore, some time-consuming operations might take longer than expected because the update has to be propagated to all nodes.

Error Recovery

The Active Cluster is designed to recover resiliently from both connection and node failures. The status of the Active Cluster and all of its members is provided in TDV Manager. If errors occur, such as a network connectivity problem or a failed node, the problem in most cases is resolved automatically. When intervention is required, Active Cluster provides a mechanism in Manager to resolve the problem and resynchronize the node.

Joining and Leaving an Active Cluster

Any TDV Server (Version 4.5 or later) with a valid license and digital certificate can join an Active Cluster. All TDV metadata on the server joining the cluster is overwritten with the TDV metadata from the cluster. Also all TDV servers in an Active Cluster must be at the same version and patch level for both TDV and Active Cluster.

If a member of the cluster leaves the cluster, the remaining members of the cluster are made aware of the member leaving and continue to operate without any need for intervention. If the member rejoins the cluster, its metadata can be merged with the cluster or replaced with the cluster metadata. Manager displays the status of the cluster so that you always know the status of all members of the cluster. A TDV Server can be added or removed at any time.

Security Implementation

Active Cluster uses the SSL (secure socket layer) communication protocol to provide secure communications. Every TDV Server has an SSL key for inter-node communications. In addition, Active Cluster uses certificate-based authentication. Every TDV Server in an Active Cluster must have a digital certificate which identifies and validates the source of the data transferred between the TDV Servers in the cluster.

All Active Cluster management operations require administrative privileges.

Installation

The installation procedure for Active Cluster is in the *TDV Installation and Upgrade Guide*.

Planning the Deployment of Active Cluster

This topic describes important considerations for implementing Active Cluster.

- [Considerations for Deploying Active Cluster, page 17](#)
- [Deployment Example, page 19](#)
- [Typical Process for Deployment, page 19](#)

Considerations for Deploying Active Cluster

Considerations and guidelines for your cluster strategy are discussed in the following:

- [Deployment Environments, page 17](#)
- [Determining the Number of Nodes in a Cluster, page 18](#)
- [Performance Considerations, page 18](#)
- [Monitor Considerations, page 18](#)
- [Caching, page 18](#)

Deployment Environments

Because of the availability and failover advantages of using a cluster, we recommend that Active Cluster be used in your QA/staging and production environments. Follow these guidelines:

- Deploy at least two clusters – one for QA/staging and one for production.
Note: we recommend *not* using clusters in your development environment because of the potential for frequent metadata changes, which can destabilize the system.
- Enterprise deployments should consider deploying multiple clusters.
- For production nodes, use 2+ CPU cores. These are minimum requirements. The actual configurations depend on such factors as the number of expected requests, number of clients, level of availability required, and so on.
- If you want to run Active Cluster with your development instances, use multiple DEV instances, but do not cluster them together.

An Active Cluster license must be installed on every cluster node. See *TDV Installation and Upgrade Guide* for how to do this.

Determining the Number of Nodes in a Cluster

For small cluster configurations where you want to ensure high availability, deploy each Active Cluster with at least three physically separate cluster nodes (separate TDV servers). Each cluster node should be deployed on a computer with two or more CPUs.

Estimate the number of users and client requests that you expect to operate simultaneously.

Finally, consider how many servers might fail simultaneously. Be sure that if the maximum number of failures occurred, the remaining server nodes can handle the work load.

Having a large number of nodes in a cluster can have an impact on performance. See [Performance Considerations, page 18](#).

Performance Considerations

In a typical development environment, where metadata is frequently changing, having too many cluster nodes can reduce performance because of the need to synchronize them. In a production environment, when the metadata rarely changes, the impact of having many nodes is less likely to affect performance.

Monitor Considerations

To avoid seeing duplicate nodes in TDV Monitor:


- Nodes in the cluster should have the Monitor Server disabled.
- Nodes in the cluster should have no settings for the Monitor Server.
- The node acting as the Monitor Server should have no data collector enabled.

Caching

Use a database cache rather than a file-system cache for query results. A database cache can be centralized and shared; a file-system cannot.

Deployment Example

The following example shows a simple three-node cluster, each with a local metadata repository.



Typical Process for Deployment

1. Install Active Cluster on each TDV Server that will be in the cluster.

Note: All TDV Servers to be in the cluster must run the same version and patch level of TDV and Active Cluster, and have the same cluster settings. For example, heartbeat time.

2. License Active Cluster on each TDV server in the cluster.
3. Create a cluster on the first TDV server.
4. Add the other TDV servers to the cluster.
5. Monitor Active Cluster status using Manager.

If your deployment process involves multiple environments for development, staging, and production, installing Active Cluster on two nodes in the staging environment should be sufficient for testing an n-node cluster.

Working with Active Cluster

This topic describes how to create a cluster, add and remove nodes, view cluster status, and delete a cluster:

- [About Working with Active Cluster, page 21](#)
- [Creating a New Active Cluster, page 25](#)
- [Adding a TDV Server to an Active Cluster, page 27](#)
- [Renaming an Active Cluster, page 32](#)
- [Removing a TDV Server from an Active Cluster, page 34](#)
- [Propagate a Domain Between Clusters, page 36](#)
- [Delete an Active Cluster, page 36](#)
- [Viewing the Status of an Active Cluster, page 36](#)
- [Configuring an Active Cluster, page 40](#)
- [Changing the Port of a Cluster Node, page 44](#)

About Working with Active Cluster

You can work with Active Cluster using:

- Manager (Web interface)
- Command-line
- Studio

Generally, you use Manager to perform most Active Cluster tasks. In fact, we recommend that you keep the CLUSTER MANAGEMENT page in Manager available to monitor the health of the cluster. See [Viewing the Status of an Active Cluster, page 36](#) for more information.

The following table summarizes the tasks and user interfaces you use to perform them.

Task	Manager	Studio	Command-Line Interface
License the Active Cluster product	YES	YES	

Task	Manager	Studio	Command-Line Interface
View Active Cluster status	YES	YES	
Create a new Active Cluster	YES	YES	YES
Add a TDV Server node	YES	YES	YES
Remove a TDV Server node	YES	YES	YES
Repair a TDV Server node	YES	YES	YES
Rename an Active Cluster	YES	YES	YES
View the SYS_CLUSTER table		YES	
Access the Web Services for Active Cluster		YES	
View and modify Active Cluster configuration information		YES	YES

Note: You can use the Studio Web Services cluster-related operations in Data Services/Web Services/system/admin/server/operations. See the *TDV User Guide* for information.

See these topics for other aspects of working with Active Cluster:

- [About Manager with Active Cluster, page 22](#)
- [About the cluster_util Program, page 23](#)
- [About Studio with Active Cluster, page 23](#)
- [Privileges Required for Working with Active Cluster, page 23](#)
- [Triggers and Active Cluster, page 24](#)
- [prePropagation of Custom Adapters in an Active Cluster, page 25](#)

About Manager with Active Cluster

You use the Manager Web interface to:

- License the Active Cluster product.

- Create a new Active Cluster.
- View the status of an Active Cluster.
- Add a TDV Server node to a cluster.
- Rename an Active Cluster.
- Remove an Active Cluster node.
- Repair an Active Cluster.

Note: When you use Manager to manage an Active Cluster, you must connect directly to a TDV Server node in the cluster. Do not use a load balancer. Also, if the system has a load balancer and you add or remove a server from an Active Cluster, make sure that the load balancer is synchronized with the Active Cluster.

About the cluster_util Program

You can use the cluster_util script to create and manage an Active Cluster. However, you can obtain more information about configuration and status if you use the Manager. This script does not manage licensing.

See [The cluster_util Command-Line Program, page 57](#), for a description of the complete syntax of this utility and its subcommands.

About Studio with Active Cluster

You use Studio to:

- View the TDV database system table. See [Active Cluster SYS_CLUSTER System Table, page 63](#).
- Access the Web Services for Active Cluster.
- View the configuration parameter settings for Active Cluster. See [Configuring an Active Cluster, page 40](#), for more information.

Note: It is not advisable to change configuration settings without assistance from the TDV support staff.

You cannot connect an Active Cluster to a load balancer using Studio.

Privileges Required for Working with Active Cluster

All actions you can perform to create or modify an Active Cluster require administrative privileges. Specifically, you must be a member of the “admin” group on the TDV server on which any of these operations is performed:

- Create cluster
- Join cluster
- Join cluster on remote server
- Remove server from cluster
- Set cluster name
- Repair cluster

See the *TDV Administration Guide* for information about the built-in admin group and how to manage group membership.

Triggers and Active Cluster

Active Cluster handles triggers automatically. The timekeeper controls triggers, ensuring that a trigger occurs only as configured.

By default, when a trigger is executed, it is executed on each node in the cluster. However, for timer event triggers, the default is to execute the trigger once per cluster.

Active Cluster handles triggers in this way:

1. When a trigger fires in a TDV instance, the trigger manager checks to see if the trigger should be executed locally. Certain triggers, such as cache refreshes for the file cache, are always executed locally.
2. If a trigger is not local, the TDV instance that invokes the trigger is determined by the timekeeper based on a weighted round-robin policy. Each TDV instance in the cluster participates with a fixed weight of one (1) except for the timekeeper instance, whose weight is determined by the Cluster trigger distribution Weight of time keeper configuration setting. (The default is 1.)

You can set a timer event trigger to be invoked on every node in the cluster. See the *TDV User Guide* for information about how to set up triggers. Clear the “Only once per cluster” option to change the default trigger behavior.

You can also use these cluster system events as trigger conditions:

- ClusterServerConnected
- ClusterServerDisconnected
- ClusterServerShunned

See [Configuring an Active Cluster, page 40](#), for more information about configuring Active Cluster.

prePropagation of Custom Adapters in an Active Cluster

When you create an adapter in Studio, the adapter definition immediately propagates across the cluster. In particular, the adapter directory ('<TDV_install_dir>/conf/adapters/<custom or system>/<adapter_name>') is created on all nodes. However, any files that are then added to the adapter directory are NOT propagated automatically.

You can propagate an adapter in a cluster by:

- Creating the adapter on a development system
- Populating the adapter directory with all needed files
- Creating a data source using the adapter
- Exporting the data source (as a CAR file)

By importing the CAR file with the data source into a cluster node, the adapter and all the files in the adapter directory are propagated across the cluster.

Creating a New Active Cluster

After you have designed your Active Cluster implementation, you are ready to create a cluster. The general process is to choose the TDV Server that is to be the first node in the cluster, and create the cluster on this computer. You then can add other servers to the cluster. All TDV metadata and configuration settings (except computer identity information) on the first TDV Server in the cluster are adopted by the TDV Servers you add to the cluster, except for customized ldap.properties or data source capabilities files. See TDV File Customizations in the *TDV Installation and Upgrade Guide*.

You must install TDV and Active Cluster prior to creating a cluster. You must also license Active Cluster on the computer on which you are creating the cluster. See *TDV Installation and Upgrade Guide*.

Note: The version of Active Cluster must match the version of TDV, and TDV Server can belong to only one Active Cluster at a time.

You can create a new Active Cluster using either the Manager Web interface or the command-line interface. At a minimum, you must have administrative privileges (MODIFY_ALL_CONFIG) and know the computer hostname, username, and password.


To create a new Active Cluster using Manager

- 1. Run TDV Server on the computer on which you are creating a new Active Cluster.
- 2. Log in to Studio and choose Launch Manager (Web) from the Administration menu.

Manager opens in your web browser.


- 3. Log in to Manager.
- 4. From the CONFIGURATION tab, choose Cluster.

Because this TDV Server does not belong to a cluster, two buttons are displayed: Create a cluster and Join a cluster



- 5. Click Create a cluster.
- 6. Enter a new cluster name and click OK.

Manager displays the CLUSTER MANAGEMENT page with this TDV Server as its only member. See [Viewing the Status of an Active Cluster, page 36](#), for more information.



To create a new Active Cluster using the command-line interface

- 1. Run TDV Server.
- 2. Open a DOS window or UNIX terminal and navigate to the bin directory of the TDV installation directory, which by default is:
`<TDV_install_dir>\bin`
- 3. Enter the command:

```
cluster_util.bat (Windows) | cluster_util.sh(UNIX) -server
<hostname> [-port <port>]
-user <username> -password <password> [-domain <domain>]
-create
-clusterName <clustername>
[-debug]
```

Replace the information in the angle brackets with the appropriate information for this TDV Server. Information in the square brackets is optional.

For example:

```
$ ./cluster_util.sh -server localhost -user admin -password admin
-create -clusterName DevCluster
```

This command creates a cluster named DevCluster.

The cluster_util subcommands for creating a cluster are described in [The -create Subcommand](#), page 60.

See [Viewing the Status of an Active Cluster](#), page 36 for how to view the cluster and its only member.

Adding a TDV Server to an Active Cluster

After you have created a cluster, you can add TDV Servers to it. The TDV Server you add to an existing cluster becomes an exact copy of the Active Cluster server it is joining except for computer identity information. That is, all metadata, configuration, and other TDV information on the TDV Server joining the cluster is overwritten with the TDV Server information from the cluster.

When joining a TDV Server to a cluster, if you have Copy Repository Database For Cluster Join enabled in the server configuration, then the TDV Server will also adopt the administration credentials used by existing TDV Servers within the cluster.

Prior to joining a cluster, you must have installed TDV 4.5 or later and Active Cluster on the TDV Server you are adding. All members of a cluster must run the exact same version of TDV and must have the same heartbeat setting. See *TDV Installation and Upgrade Guide* for installation information. See [The Timekeeper Cluster Node Fails or Is Taken Offline](#), page 53 for information about setting the heartbeat interval.

You must also license Active Cluster on every TDV Server cluster node. See *TDV Administration Guide* for more information about licensing.

Consider the following restrictions:

- A TDV Server can belong to only one Active Cluster at a time.
- If you customized either the `ldap.properties` file or the data source capabilities file on the TDV Server in the cluster you are joining, you must manually copy these files to all computers joining the cluster. These files are not automatically synchronized by Active Cluster. See *TDV Installation and Upgrade Guide* for more information.
- If you are using a load balancer (recommended) and you add a TDV Server to a cluster, you must also add the TDV Server computer to the load balancer. Active Cluster is not aware of whether you use a load balancer with a cluster.
- The cluster join process involves adding, updating, and deleting metadata on the joining node. This can be time-consuming, especially if the joining node or the cluster itself has much metadata. See [Preparing to Join an Active Cluster, page 28](#) for tips on how to save your metadata and expedite the join process.

This section contains the following topics:

- [Preparing to Join an Active Cluster, page 28](#)
- [Joining an Active Cluster, page 29](#)

Preparing to Join an Active Cluster

To expedite the join process, make sure that the TDV Server that is to join the cluster contains minimal metadata beforehand. Because all existing resources on the joining node are overwritten with the cluster data at join time, it is best to prepare the joining node with a “clean” repository. For example, the repository has only what is present the TDV installer is run.

If the joining node has resources that you want to save and potentially restore, be sure to follow the process prior to joining the cluster.

To prepare a TDV Server with existing resources prior to joining a cluster

1. Run a `backup_export` of the system, and validate the CAR file by importing it into another system.

Validation is important, so you are sure to have a full backup file in case you want to restore data later.
2. Use the `repo_util.sh -DropSchema` command.
3. Restart the server to recreate the schema automatically.
4. Run `backup_import` of the CAR file that was previously been exported, if you want to restore the original data or resources.

See the *TDV Administration Guide* for more information about using the `backup_export`, `backup_import`, and `repo_util` utilities.

Joining an Active Cluster

You can use the Manager Web interface or the command-line interface to add a TDV Server to a cluster. At a minimum, you must know the computer hostname, username, password, and port information for both the TDV server you are adding and any one of the servers in the cluster you are joining. You must have administrative privileges (`MODIFY_ALL_CONFIG`) on both the local server and the remote server.

Adding a node to the cluster can be time-consuming especially if the cluster has large metadata size. TDV Server has two mechanisms to handle metadata replication to the joining node. The default mechanism uses the same technique as doing a full backup export and import of an archive file, which includes all clustered artifacts referenced in "Metadata Repository." An alternate mechanism exists to speed up the performance of metadata replication by copying the TDV filesystem and repository artifacts. This mechanism can be enabled using the configuration: `Server > Cluster > "Copy Repository Database For Cluster Join"`. Note that this setting needs to be enabled when you first create a cluster.

This option comes with following limitations.

- The joining node will be restarted after the cluster join.

The administrator joining a node to the cluster will see an informational message confirming that the joining node is going to be restarted after the join. The administrator would need to re-login to the joining node after it has successfully restarted.

- The admin and monitor passwords will be overwritten by the ones from which the node replicates its data.

It is generally recommended to maintain the same admin and monitor passwords for a cluster and this is strictly enforced when the "Copy Repository Database For Cluster Join" configuration is enabled.

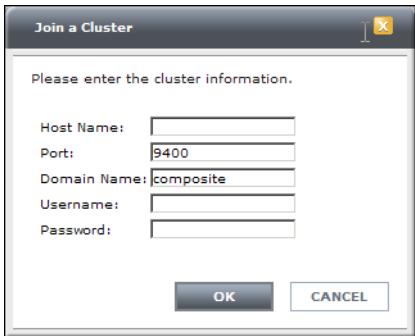
To join an Active Cluster using Manager

1. Run TDV Server on the computer which you are adding to the Active Cluster.
2. In Studio, choose Launch Manager (Web) from the Administration menu.
Manager opens in your web browser.
3. Log in to Manager.
4. From the CONFIGURATION tab, choose Cluster.

Because this TDV Server does not yet belong to a cluster, two buttons are displayed: Create a cluster and Join a cluster.

- 5. Click Join a cluster.

Manager opens a dialog box for you to specify information for a member of the cluster:



- 6. Enter the information for any one of the TDV Servers that are in the cluster you are joining, and click **OK**.

Field	Value
Hostname	The hostname of any TDV Server that is already a member of the cluster.
Port	The port ID of the TDV Server that is in the cluster. By default, the port ID is filled in.
Domain Name	The name of the domain to which the cluster server belongs. By default, the domain for this TDV server is filled in.
Username	The name of an administrative user who can access the cluster server.
Password	The password for the specified username.

Note: Be sure to enter the physical TDV Server address information—not the load balancer address information.

Active Cluster overwrites the metadata and configuration settings in the current TDV instance with the metadata and settings of the cluster. Depending on the volume of metadata involved and network latency between nodes, this operation may take several minutes to complete.

You can use Manager on another node to see cluster status of joining node during a long-running join process. You can also monitor the `cs_server_metadata.log` in `<TDV_install_dir>\logs\` to see the changes as

they are applied. By examining the metadata log file, you can see what objects are being written to the repository, and track how the joining process is going.

Note: During the join operation, you cannot run Studio or perform any activities that use the joining TDV Server in Manager.

If the join is successful, a success message is displayed.

Manager then displays the CLUSTER MANAGEMENT page with a row for every TDV Server that belongs to the cluster, including a row for the TDV server you just added. See [Viewing the Status of an Active Cluster, page 36](#) for the meaning of the columns and the status information.

CLUSTER MEMBERSHIP

Here you can manage cluster membership.

Cluster:

Server Time: 4/26/2010 11:33:44 am

Successfully refreshed at: 4/26/2010 11:31:23 am

Refresh Rate:

	Server	Status	Host	Port	Timekeeper	Connect Time	Last Co
	jaudette.composite.com-9410--744357719	CONNECTED_READY	jaudette.composite.com	9410	Y	4/26/10 6:31 PM	4/26/10 6:31 PM
	jaudette.composite.com-9400--349377532	OPERATIONAL [36]	jaudette.composite.com	9400	N		

Rename Cluster Remove Selected Cluster Node Repair This Node

All other members of the cluster immediately become aware of the new member, and it is automatically synchronized when updates occur.

The Status column indicates the status of the node. The cluster node you are currently connected to in the browser is indicated by an OPERATIONAL status code; the other operational nodes have a status of CONNECTED_READY. See [Viewing the Status of an Active Cluster, page 36](#) for a complete description of the columns in this table and the other possible Status codes.

You can check the server logs for Active Cluster activity. These logs are in the <TDV_install_dir>\logs directory. All Active Cluster log messages are written to the cs_cluster.log in the cluster folder under the logs directory of each node. For the joining node, metadata changes (deletions and additions) are signed in cs_server_metadata.log.

To add a TDV Server to an existing cluster the command-line interface

1. Run TDV Server on the computer which you are adding to the Active Cluster.
2. Open a DOS window or UNIX terminal and navigate to the bin directory of the TDV installation directory, which by default is:
<TDV_install_dir>\bin

3. Enter the command:

```
cluster_util.bat (Windows) | cluster_util.sh (UNIX)
-server <target server hostname> [-port <target server port>]
-user <username> -password <password> [-domain <domain>]
-join
-memberServer <cluster server hostname> [-memberPort <cluster
server port>]
-memberUser <username> -memberPassword <password> [-memberDomain
-domain <domain>]
[-debug]
```

Identify the target server that is joining the cluster and one of the remote servers that is a member of the cluster. You must have administrative credentials on both servers. Information in the angle brackets is required. Information in the square brackets is optional.

At a minimum, you must specify the name of the new cluster server node (the target server), its username and password, and the server name, username, and password for a remote server that is already a member of the cluster. For example:

```
$ ./cluster_util.sh -server localhost -user userd1 -password
password1
-join -memberServer Bostonhost -memberUser admin -memberPassword
admin
```

This command joins a computer named localhost to a cluster node named Bostonhost.

After the server joins, the server immediately becomes a replicate of all other servers in the cluster.

Note: When you add a TDV Server to an Active Cluster, any existing metadata and the configuration settings (except computer identity information) are overwritten with the metadata and configuration settings from the cluster. That is, all existing TDV Server information on the computer you are adding to the Active Cluster is lost.

The cluster_util subcommands for joining a cluster are described in [The -join Subcommand](#), page 61.

Renaming an Active Cluster

The name of an Active Cluster appears in the Manager on the CLUSTER MANAGEMENT page. You can rename an Active Cluster at any time. The new name is automatically propagated to all other cluster members.

You can rename an Active Cluster using the Manager Web interface or the command-line interface.

To rename an Active Cluster using Manager

1. Run TDV Server on one of the servers that belongs to the cluster you want to rename.
2. Log in to Studio and choose Launch Manager (Web) from the Administration menu.

Manager opens in your web browser.

3. Log in to Manager.
4. From the CONFIGURATION tab, choose Cluster.
The members of the Active Cluster are displayed along with their status.
5. Click Rename Cluster.

Manager opens a dialog box for you to specify a new name.

6. Enter a new cluster name and click **OK**.

Manager displays the new cluster name of the cluster status table and propagates the name to all other members of the cluster.

To rename an Active Cluster using the command-line interface

1. Run TDV Server on one of the servers that belongs to the cluster you want to rename.
2. Open a DOS window or UNIX terminal and navigate to the bin directory of the TDV installation directory, which by default is:

<TDV_install_dir>\bin

3. Enter the command:
`cluster_util.bat (Windows) | cluster_util.sh (UNIX) -server`
`<cluster server hostname>`
`[-port <cluster server port>]`
`-user <username> -password <password> [-domain <domain>]`
`-setClusterName`
`-clusterName <clustername>`
`[-debug]`

Replace the information in the angle brackets with the appropriate information for a TDV Server in the cluster. Information in the square brackets is optional.

At a minimum, you must specify the name of a cluster server node, its username and password, and the new name for the cluster. For example:

```
$ ./cluster_util.sh -server localhost -user admin -password admin
-setClusterName -clusterName ProdCluster
```

This command assigns the new name ProdCluster to the cluster that the server localhost belongs to.

The cluster_util subcommand for renaming a cluster are described in [The -setClusterName Subcommand](#), page 62.

Removing a TDV Server from an Active Cluster

You can remove a TDV Server from an Active Cluster:

- From the TDV Server node that you are removing.
- From any other TDV Server in the cluster.

In both cases, all remaining cluster nodes are automatically notified that the TDV Server has been removed.

When you remove a TDV Server from an Active Cluster, it becomes a stand-alone server that initially has all of the metadata and configuration settings of the cluster.

Note: You can use the removed cluster node to create a new cluster. For example, if you have developed and tested a particular TDV Server implementation and are ready to put it into a production environment.

If you are using a load balancer and you remove a TDV Server from a cluster, the load balancer still sends work to the removed server. You need to remove the TDV Server from load balancer's configuration to make sure that requests intended for the cluster are known to all servers.

You can remove a TDV Server cluster node using the Manager Web interface or the command-line interface. You must know the computer name, username, password, and port information.

To remove a TDV Server from a cluster using Manager

1. Run TDV Server on the computer that you are removing from the Active Cluster.
2. In Studio, choose Launch Manager (Web) from the Administration menu.
Manager opens in your web browser.

3. Log in to Manager.
4. From the CONFIGURATION tab, choose Cluster.
The members of the Active Cluster are displayed along with their status.
5. Click the radio button next to the name of the TDV Server node you want to remove.
6. Click Remove Selected Cluster Node.
The TDV Server becomes a stand-alone server that retains all data that it had as part of the cluster.

To remove a TDV Server from a cluster using the command-line interface:

1. Run TDV Server on the computer that you are removing from the Active Cluster.
2. Open a DOS window or UNIX terminal and navigate to the bin directory of the TDV installation.
3. Enter the command:
`cluster_util.bat (Windows) | cluster_util.sh (UNIX)`
`-server <cluster server hostname> [-port <cluster server port>]`
`-user <username> -password <password> [-domain <domain>]`
`-remove`
`[-memberServer <target server hostname>]`
`[-memberPort <target server port>]`
`[-debug]`

Replace the information in the angle brackets with the appropriate information for a TDV Server in the cluster. Information in the square brackets is optional.

At a minimum, you must specify the name of a cluster server node you are removing and its username and password. For example:

```
$ ./cluster_util.sh -server localhost -port 9410 -user admin
-password admin
-remove
```

This command removes the cluster node with the alias of localhost.

The cluster_util subcommands for removing a cluster node are described in [The -remove Subcommand, page 62](#).

Propagate a Domain Between Clusters

An Active Directory domain created in one cluster can be propagated correctly to another cluster, but only if you set the domain to the external domain first and then clone the environment onto different cluster nodes. To propagate the `ldap.properties` file across the cluster, synchronize the nodes first and then apply changes manually.

If the domain type was changed to external domain after a node was joined to a cluster, take the node out of the cluster and then rejoin it to the cluster.

Delete an Active Cluster

To delete a cluster, you remove each TDV Server node, one at a time, as described in [Removing a TDV Server from an Active Cluster, page 34](#). Removing the last remaining TDV Server cluster node deletes the cluster.

Viewing the Status of an Active Cluster

You can see the status of all TDV Servers in the cluster using the Manager Web interface. Administrators have full access to this information. You can view this the cluster status information read-only if you have `ACCESS_TOOLS` and either `READ_ALL_CONFIG` or `READ_ALL_STATUS` privileges.

Note: Be sure to connect directly to a cluster node when using the Manager to view the cluster status. Do not connect using a load balancer.

To view the status of an Active Cluster and its members

1. Run TDV Server on a computer that belongs to the Active Cluster.
2. In Studio, choose Launch Manager (Web) from the Administration menu.
Manager opens in your web browser.
3. Log in to Manager.
4. From the CONFIGURATION tab, choose Cluster.

If an Active Cluster has been defined, the members of the cluster are displayed along with their status.

Note: To make sure that the status information is up to date, click the Refresh Now icon.

CLUSTER MEMBERSHIP

Here you can manage cluster membership.
Cluster: MyCluster

Server Time: 4/26/2010 2:08:25
Successfully refreshed at: 4/26/2010 11:42:5
Refresh Rate: Never

	Server	Status	Host	Port	Timekeeper	Connect Time	Last Contact	Idle	Connect Count	Bytes To Server	Byte
<input type="radio"/>	jaudette.composite.com-9410--744357719	OPERATIONAL [36]	jaudette.composite.com	9410	Y			0	0	0	0
<input type="radio"/>	jaudette.composite.com-9400--349377532	CONNECTED_READY	jaudette.composite.com	9400	N	4/26/10 6:42 PM	4/26/10 6:42 PM	44767	3	227156	1834

Rename Cluster

Remove Selected Cluster Node

Repair This

The TDV Server you are connected to is always listed, and can be identified by the Host and Port information.

Information about the TDV Server node to which you are connected is provided:

- Server Time—The current time.
- Successfully refreshed at—The last time the data for this server was synchronized with the cluster.
- Refresh Rate—If you want the cluster status information to refresh periodically, select the length of time between refreshes. You can instantaneously refresh the server by clicking the Refresh Now icon.

The following table describes the status information provided in the table for each TDV Server node that belongs to the Active Cluster.

Column	Description
<radio button>	Select to take an action with this TDV server node.
Server	The full display name of this server as defined by TDV.

Column	Description
Status [#] (for a local node)	<p>For the TDV server to which you are connected), displays one of these values:</p> <p>INITIALIZING_LOCAL—This TDV server node is initializing its internal state.</p> <p>INITIALIZING_REMOTE—This TDV server node is initializing its connections with other remote cluster nodes.</p> <p>OPERATIONAL—This TDV server node is fully operational and synchronized.</p> <p>BLOCKED—This TDV server node is awaiting a response from one or more server nodes for a period of time that exceeds the Block Threshold specified in Studio.</p> <p>LEAVING—This TDV server node is being removed from the cluster.</p> <p>The Message column provides additional information pertaining to the status. For OPERATIONAL nodes, the Message column lists the pending responses from remote nodes this node is waiting on.</p> <p>[#]—The change level for this node. This value indicates the level at which changes have been captured in the metadata repository. In steady state, all nodes in a cluster should be operating at the same change level. However, because changes can occur on other nodes at any time, only the local node change level is known and displayed.</p>
Status (for a remote node)	<p>For a remote node, displays one of these values:</p> <p>CONNECTED—This TDV server node is connected to the cluster, although not necessarily synchronized with the rest of the cluster.</p> <p>CONNECTED_READY—This TDV server node is connected to the cluster and is synchronized and ready to receive change notifications.</p> <p>BLOCKING—This TDV server node has not responded to another server node for a period that is longer than the Block Threshold specified in Studio. See Setting a Timekeeper and Repairing Active Cluster Nodes, page 48 to repair a node in this state.</p> <p>DISCONNECTED—This TDV server node is not connected to the cluster.</p> <p>INVALID_LICENSE—This TDV server is attempting to join the cluster but does not have a valid license.</p>
Host	The hostname for this server.
Port	The port to which this TDV server node is connected.

Column	Description
Timekeeper	<p>Displays one of two values:</p> <p>Y—This TDV server node is the timekeeper for the cluster.</p> <p>N—This TDV server node is not the timekeeper for the cluster.</p> <p>See About the Timekeeper, page 14, for more information.</p>
Connect Time	The duration, in minutes, of the connection established between the local TDV server node and the cluster node displayed in that row. This value is updated each time the server connects to a cluster node and becomes an active member of the cluster. This value defaults to GMT (Greenwich Mean Time), but you can change the time zone using the Web Services Interface configuration parameters.
Last Contact	The last time that a message was received across the cluster network connection by the local TDV Server node. This value defaults to GMT (Greenwich Mean Time), but can be changed.
Idle	Number of milliseconds between the Last Contact time and now.
Connect Count	Number of successful connections to the cluster by the local TDV server node.
Bytes To Server	Number of bytes sent to the local TDV server node since it started.
Bytes From Server	Number of bytes sent from the remote TDV server to the local TDV server node.
Message	Additional information about the activity of the TDV server. Also indicates the cluster nodes from which this node is awaiting responses.

Getting the Active Cluster and Member Names

The best way to get information about an Active Cluster and its members is using the Manager, as described in [Viewing the Status of an Active Cluster, page 36](#). However, you can get the name of an Active Cluster along with its member names using the command-line interface.

To get the active cluster and its member names

1. Run TDV Server on the computer for which you want configuration information.
2. Open a DOS window or UNIX terminal and navigate to the bin directory of the TDV installation.
3. Enter the command:

```
cluster_util.bat (Windows) | cluster_util.sh (UNIX)
-server <cluster server hostname> [-port <cluster server port>]
-user <username> -password <password> [-domain <domain>]
-getConfig
[-debug]
```

Replace the information in the angle brackets with the appropriate information for this TDV Server. Information in the square brackets is optional.

At a minimum, you must specify the name of a cluster server node and its username and password. For example:
\$./cluster_util.sh -server localhost -user admin -password admin -getConfig

This command retrieves the cluster name and the cluster member names for the cluster to which the server localhost belongs.

The cluster_util subcommand for getting cluster configuration information for a cluster is described in [The -getConfig Subcommand, page 61](#).

Configuring an Active Cluster

You can configure some aspects of an active cluster using the configuration parameters in Studio. This section describes the active cluster configuration parameters and how to reconfigure a node’s port number if necessary.

To configure an Active Cluster

- 1. Open Studio on one of the cluster nodes.
- 2. From the Administration menu, choose Configuration.

The purpose of the Active Cluster configuration parameters is described in the following table.

Cluster Configuration Parameter	Description	Default	Min
Block Threshold	Time threshold in minutes beyond which the TDV server considers itself BLOCKED if it has not received a pending response.	60 minutes	1

Cluster Configuration Parameter	Description	Default	Min
Global Change ID	<p>The current global change ID is a counter for each change that occurs; it is the value displayed in brackets next to the status in Manager. This value matches across all nodes, although it might not match until changes are propagated. The IDs can differ by a maximum of one. Differences of 2 or more indicate a problem to investigate.</p> <p>This is a read-only value that indicates the current server state.</p>	-1	N/A
Copy Repository Database for Cluster Join	<p>A mechanism used to speed up the performance of metadata replication by copying the TDV filesystem and repository artifacts. It is set to False by default and needs to be enabled when you first create a cluster.</p> <p><i>Note:</i></p> <ul style="list-style-type: none"> The joining node will be restarted after the cluster join. The administrator would need to re-login to the joining node after it has successfully restarted. The admin and monitor passwords will be overwritten by the ones from which the node replicates its data. 	False	N/A
Health Monitor Aggressive Heartbeat Interval	After any node is flagged as down, specifies how often to check for the heartbeat of other nodes, in seconds.	300 seconds	1
Health Monitor Heartbeat Table Path in TDV (Optional Feature)	<p>Health Monitor heartbeat table path in TDV. This configuration parameter must have a value specified for this optional feature to work.</p> <p>The Cluster Health Monitor Table must be created using the Info tab for the data source where you want the heartbeat data cached. Using the Browse button leads you through a set of screens that you can use to have TDV create the table.</p>	<null>	N/A
Health Monitor Heartbeat Interval	<p>Number of seconds between heartbeats.</p> <p>Frequency that the new Health Monitor table is updated with status.</p>	600 seconds	1

Cluster Configuration Parameter	Description	Default	Min
Health Monitor Heartbeat Read Interval	Number of seconds to wait after sending a heartbeat before reading heartbeats of other nodes in cluster.	30 seconds	1
Health Monitor Tolerance Interval	Number of seconds to wait after a node is unresponsive before flagging it as down.	600 seconds	1
Cluster Heartbeat Interval (On Server Restart)	Number of seconds between heartbeat messages the server sends out to other cluster members. This is a read-only value that indicates the current cluster heartbeat interval. To change this value, use the Cluster Heartbeat Interval configuration parameter.	5 seconds	1
Cluster Heartbeat Interval	Number of seconds between heartbeat messages Server sends out to other cluster members. Changing this value has no effect until the next server restart.	5 seconds	1 second
Cluster Logging Detail Level	Can be set to: OFF—No detailed logging is enabled. MEMBER—Log member and timekeeper changes only. METADATA—Log metadata changes in addition to member and timekeeper changes. ALL—Log all cluster communications. See the cs_server.log and cs_server_metadata.log in the <TDV_install_dir>\logs directory. See Configuring Logging for Active Cluster, page 43 for information about other logging settings.	OFF	N/A
Repository Retention Period	The minimum number of hours the server is to keep a metadata object after it is deleted. This is needed only if some servers in the cluster are not up all the time, or the network among the cluster members are intermittent. This setting is ignored if all servers in the cluster are running and communicating with each other.	72 hours	1

Cluster Configuration Parameter	Description	Default	Min
Connection Retry Interval	The time interval in seconds between attempts to contact a nonresponding remote cluster node.	5 seconds	0
Initial Response Timeout	The number of seconds the server is to wait for an initial response while establishing a connection to a remote cluster node.	5 seconds	0
Weight of time keeper	Set this if timekeeper needs special treatment in weighted round robin trigger distribution. For most use cases, this value does not need to be set.	1	0

Configuring Logging for Active Cluster

Logging of system events is controlled by Studio configuration parameters. By default, whenever an Active Cluster node is connected, disconnected, or shunned (not allowed into the cluster), the event is logged as specified in the Cluster Event configuration parameters.

To configure cluster events for logging

- 1. In Studio, choose Administration > Configuration from the main menu.
- 2. Adjust the Cluster Event configuration settings as desired.

Configuration Parameter	What It Specifies
Enable All Events	Whether to log Active Cluster events according to the settings specified by the parameters. True (default)—Enable all events in this category. False—Disable all events in this category.

Configuration Parameter	What It Specifies
Server Connected	<p>How server-connect events are logged. You can specify multiple choices, separated by commas (except for ALL and NONE); for example, DB,LOG,SNMP.</p> <p>The event filters are:</p> <p>DB—Event sent to database only.</p> <p>LOG—Event sent to cs_server_events.log file only. This file resides in the <TDV_install_dir>\logs directory.</p> <p>SNMP—Event sent to SNMP processor only.</p> <p>CUSTOM—Event sent to custom event handler only.</p> <p>ALL—Event sent to database, log, SNMP processor, and custom event handler.</p> <p>NONE—Event ignored.</p>
Server Disconnected	<p>How server-disconnect events are logged. Valid filters are the same as the Server Connected options.</p>
Server Joined	<p>How server-join events are logged. Valid filters are the same as the Server Connected options.</p>
Server Shunned	<p>How server-shunned events are logged. Valid filters are the same as the Server Connected options.</p>

Changing the Port of a Cluster Node

If you need to change the port number of a cluster node, you can do this. However, because the port number is part of the name of the server instance in the cluster, changing the port number causes it to not be recognized as a cluster member by the rest of the cluster members. Follow the procedure to change the port number of a cluster node that you want to be a part of the same cluster.

Changing the HTTP base port value also changes the value of all derived ports after the next TDV restart (with the exception of the Repository and Cache database ports, which will remain the same).

To change the port number of a cluster node

1. In Manager, remove the TDV Server node from the cluster.
2. Change the TDV Server port number in Studio:
 - a. Open Studio on one of the cluster nodes.
 - b. From the main menu, choose Administration > Configuration.
 - c. Locate Port (Current) to display the current port setting for this TDV Server.
 - d. Change the port number.
 - e. Click **OK**.
3. Shut down the TDV Server.
4. Restart the TDV Server so that the port change takes effect.
5. In Manager, rejoin the cluster.

Note: In the Port configuration panel, the Derived TDV Ports shows that HTTP base port 7 is reserved for Active Cluster/JGroups. Active Cluster uses the JGroups Java library to track the nodes that are up and down. Heartbeat messages are sent on a different port (the derived port) than the port used for metadata.

Handling Active Cluster Errors

This topic describes how to recover from errors in an Active Cluster.

- [Overview of Error Handling, page 47](#)
- [Setting a Timekeeper and Repairing Active Cluster Nodes, page 48](#)
- [Troubleshooting, page 50](#)
- [Other Best Practices for Active Cluster, page 55](#)

Overview of Error Handling

Active Cluster is resilient after connection and node failures; most of the time such failures are resolved automatically:

- If a node goes down, the rest of the group continues to operate as a single cluster. After the node is restored, it automatically reconnects with the rest of the group.
- If a connection failure occurs and the cluster becomes partitioned into subgroups, each subgroup operates as a separate cluster with its own timekeeper. After the connection is restored, the subgroups automatically merge to form a single group with a single timekeeper.

If you do have cluster problems, such as nodes with frequent DISCONNECTED, BLOCKED, or BLOCKING status, refer to [Troubleshooting, page 50](#).

Get Cluster Status

You can get an overall picture of the activity of a cluster on the CLUSTER MANAGEMENT page in the Manager. Nodes that are off-line or down for any reason have the status of DISCONNECTED. See [Viewing the Status of an Active Cluster, page 36](#), for more information.

The cluster status view depends on which node you are viewing in the browser (the “local node”). The local node is the one with a number in brackets in the Status column (second from the left); for example, “[36]”.

Only a local node can have the status of OPERATIONAL; all remote nodes that are operational have the status of CONNECTED_READY.

When troubleshooting, it is a good idea to log in to each node. For example, a node can be blocking other nodes, but its status is OPERATIONAL. In this case, other nodes would have status of BLOCKED, meaning the OPERATIONAL node is actually blocking the other nodes and needs repair.

Setting a Timekeeper and Repairing Active Cluster Nodes

Having an effective way to recover nodes without restarting servers is essential. Recovering cluster nodes to a working state without restarting the TDV server helps maintain stable cluster environments. Regrouping the cluster attempts to establish connections with all the nodes in the cluster and set the timekeeper node. This can help if the:

- Timekeeper is missing
- Nodes are in different states such as CONNECTED, OPERATIONAL, and DISCONNECTED.
- Cluster synchronization does not work.

You can attempt to regroup an Active Cluster node. When you regroup a node, Active Cluster resets all cluster connections for the node, and then the node rejoins the cluster and resynchronizes its metadata. Reachable nodes are regrouped to the cluster and have their connections recycled. Nodes that are shut down or unreachable are skipped.

If the node contains metadata changes that are incompatible with the rest of the cluster, it is not allowed to rejoin.

- [Regroup Using Manager, page 48](#)
- [Regroup Using the cluster_util Script, page 49](#)
- [Regroup Using the API, page 50](#)

Regroup Using Manager

To regroup using Manager

1. Run Studio and choose Launch Manager (Web) from the Administration menu.

Manager opens in your web browser.

2. Log in to Manager.
3. From the CONFIGURATION tab, choose Cluster.

The members of the Active Cluster are displayed along with their status.

4. In the Manager of the node, identify the cluster nodes that need to rejoin the cluster.
5. Log on to the Manager of the node that needs repair.
6. Click **Regroup Cluster**.

The button always operates on the node the browser is connected to. If the node that you are logged in to from manager is not currently the timekeeper, using this button will make it the timekeeper node.

Manager displays a Confirmation dialog.

7. Click **OK** to repair this node.

Active Cluster resets all cluster connections for this node, rejoins the cluster, and resynchronizes its metadata with the cluster.

8. Click **OK**.

If there are metadata conflicts, this node is removed from the cluster.

See [Troubleshooting, page 50](#) for more information about how to resolve cluster issues.

Regroup Using the cluster_util Script

The cluster_util script can issue a regroup command to all nodes that are reachable and recycle their connections while setting the timekeeper. This does not require a cluster reboot.

The cluster_util command-line parameters for repairing a cluster node are described in [Using the cluster_util Program, page 57](#).

To re-group using the cluster_util

1. From the computer where TDV is installed, open a DOS window or UNIX terminal.
2. Navigate to the <TDV_install_dir>/bin directory.
3. Run the cluster_util script using .bat for Windows or .sh for UNIX. For example:

```
./cluster_util.sh
-regroup
-user <username>
-password <password>
-server <cluster server hostname>
[-port <cluster server port>]
```

Replace the information in the angle brackets with the appropriate information for a TDV server in the cluster. Information in square brackets is optional. The node indicated as the <cluster server hostname> becomes the timekeeper node if it is not already the timekeeper node.

For Example

```
./cluster_util.sh -regroup -user admin -password admin -server
localhost -port 9400
```

Regroup Using the API

The cluster regroup API can be run as a web service or from the command line. This API issues a regroup command to all nodes that are reachable and recycles their connections while setting the timekeeper to the node where the command was issued. This does not require a cluster reboot.

From Studio you can navigate to the Web Service:

```
/services/webservices/system/admin/server/operations/repairCluster
```

When using the API for regrouping, specify REGROUP as the option.

```
<server:repairClusterxmlns:server="http://www.compositesw.com/services/system/admin/server">
  <server:option>REGROUP</server:option>
</server:repairCluster>
```

Troubleshooting

The following sections list some potential issues and suggests actions to resolve them.

- [Issues with Setting Up a Cluster, page 51](#)
- [A Cluster Node Is DISCONNECTED, page 51](#)
- [A Cluster Node Is Frequently DISCONNECTED, page 53](#)
- [The Timekeeper Cluster Node Fails or Is Taken Offline, page 53](#)
- [A Cluster Node Is Frequently BLOCKED, page 54](#)
- [Cannot start CMS after Backing Up a Cluster Node, page 54](#)
- [Other Best Practices for Active Cluster, page 55](#)

Issues with Setting Up a Cluster

Common issues that can arise when setting up a cluster are:

- **Connectivity**—Make sure that all candidate cluster nodes can communicate with each other.
- **Hostnames**—Make sure that the cluster node hostname is mapped to its actual IP address.
- **Version Compatibility**—All cluster nodes must be running the same version and patch level of TDV and Active Cluster.
- **Licensing**—All cluster nodes must be licensed prior to joining a cluster and all must have identical licensing types.

See *TDV Installation and Upgrade Guide* for details about the requirements for Active Cluster.

A Cluster Node Is DISCONNECTED

A DISCONNECTED node might require intervention:

- [A Cluster Node Goes Offline, page 51](#)
- [A Cluster Node Is Busy and Does Not Respond, page 51](#)
- [Nodes Have Been Partitioned into Subgroups, page 52](#)
- [A Node Has Been Evicted Due to Metadata Conflicts, page 52](#)

A Cluster Node Goes Offline

If a cluster node is rebooted, disconnected, or taken offline, its status changes to DISCONNECTED from the view of the remaining nodes. When the cluster node comes back up, it automatically rejoins and is resynchronized. However, the cluster node cannot rejoin the cluster if while the node was disconnected the node's metadata **and** the cluster metadata were modified. To rejoin the node to the cluster, follow the procedure in [Adding a TDV Server to an Active Cluster, page 27](#).

A Cluster Node Is Busy and Does Not Respond

Each node generates a heartbeat that alerts the cluster on base port plus seven (for example, 9407) that the node is still connected. If a node fails to generate a heartbeat within a designated period, the node is temporarily disconnected from the cluster. When the busy node next sends its heartbeat, the other nodes reject it,

prompting the busy node to reset its cluster connections and attempt to resynchronize and push its changes to the cluster. However, if metadata changes are in conflict, the node is removed from the cluster. See [A Node Has Been Evicted Due to Metadata Conflicts, page 52](#) for more information.

The cluster node can rejoin the cluster by following the procedure in [Adding a TDV Server to an Active Cluster, page 27](#). However, because joining a cluster wipes all metadata in the joining node, you might need to make your changes again if you want them to be reflected in the cluster.

Nodes Have Been Partitioned into Subgroups

If connection failures occur, the network topology and configuration might cause the cluster to be partitioned into subgroups. For example, if some nodes are connected to the cluster through a common, failed router, those nodes could become a cluster subgroup. The DISCONNECTED status of the cluster nodes in Manager can help to troubleshoot this type of event.

For example, in a five-node cluster with nodes A through E, if D and E are simultaneously disconnected, two subgroups might be formed, each with its own timekeeper: A, B, and C in one, and D and E in the other. In Manager, all cluster nodes in both groups would be visible. The status of the subgroup nodes D and E would appear as DISCONNECTED from nodes A, B, and C. The status of nodes A, B, and C would appear as DISCONNECTED from nodes D and E.

During partitioning, changes can be made in each cluster group. When connections are successfully re-established, the two subgroups are automatically merged and the metadata synchronized, if there are no metadata conflicts. The original timekeeper would again become the timekeeper for the merged cluster.

See [A Node Has Been Evicted Due to Metadata Conflicts, page 52](#) for more information.

A Node Has Been Evicted Due to Metadata Conflicts

Situations can occur where metadata changes are in conflict. For example:

- Nodes are out of sync because the cluster was partitioned, and metadata changes happened in both partitions.

In such a case, the partition with the original timekeeper prevails and the nodes belonging to the other partitions are removed from the cluster. (A partition can contain just one node.)

- Nodes are out of sync because each is modified in the absence of the other.

The presence of a third active node in these scenarios would prevent the conflict from happening, because the third node would propagate the changes to.

Typically, the node that is out of sync is automatically evicted from the cluster, and all sessions are terminated. If a node is automatically evicted, review the server and cluster logs to find the cause and resolve it.

A Cluster Node Is Frequently DISCONNECTED

If a node frequently has a status of DISCONNECTED, you can try increasing the value of the Cluster Heartbeat Interval (On Server Restart) configuration parameter to reduce the likelihood of it being disconnected, and thereby improve its performance.

To adjust the Cluster Heartbeat Interval (On Server Restart) parameter

1. Verify that all servers in the cluster are running.
2. With Studio connected to any TDV server in the cluster, choose Administration > Configuration.
3. Navigate to TDV Server > Configuration > Cluster.
4. Adjust the value of Cluster Heartbeat Interval (On Server Restart).
5. Shut down all TDV servers.
6. Restart all TDV servers.

The Timekeeper Cluster Node Fails or Is Taken Offline

If the timekeeper cluster node becomes unavailable for any reason, another TDV Server assumes timekeeper responsibilities automatically. You can see which cluster node is the timekeeper in the Timekeeper column on the CLUSTER MANAGEMENT page in Manager.

The timekeeper role assignment can change based on circumstances:

- If the original timekeeper was disconnected due to a failure but not shut down, it resumes the timekeeper role after it successfully reconnects and resynchronizes. If the cluster is partitioned into subgroups as a result of the failure, each subgroup gets its own timekeeper. See [Nodes Have Been Partitioned into Subgroups, page 52](#). The original timekeeper remains timekeeper for its subgroup and resumes the timekeeper role for the entire cluster when the subgroups are merged.
- If the original timekeeper node is being repaired, the next node that joined the cluster becomes the timekeeper. The new timekeeper keeps its role as timekeeper even after the old timekeeper node is repaired.

- If the original timekeeper fails to send a heartbeat to the cluster group, the next node that joined the cluster becomes the new timekeeper.

A Cluster Node Is Frequently BLOCKED

The cluster node to which you are connected has the status of BLOCKED if it has been awaiting a response from another node for more than the designated blocking period.

The node being waited upon has the status of BLOCKING. The BLOCKING node can be repaired using the process described in [Setting a Timekeeper and Repairing Active Cluster Nodes, page 48](#). However, if this occurs frequently and is caused by normal operations, you might want to adjust the designated blocking period, which is the maximum amount of time required for normal operations to complete.

To adjust the Block Threshold configuration parameter

1. In Studio, choose Administration > Configuration.
2. Adjust the value of the Block Threshold configuration parameter.

Cannot start CMS after Backing Up a Cluster Node

CMS operations, such as install CENTRAL mode server, start CMS, might not start as expected on the backup nodes of your Active Cluster environment. You might need to change the TDV configuration settings so that the CMS services can start automatically.

To automatically start CMS services

1. Start Active Cluster.
2. In Studio, choose Administration > Configuration.
3. Locate and set the value of the Store User Password configuration parameter to true.

This change will be synced between cluster nodes.

4. On the primary node, change the "admin" password.
5. On the backup node, change the "admin" password, the new password should be the same as the new password on the primary node.
6. Restart Active Cluster.
7. Sign in to Studio on the backup node as admin/<new password>.

Other Best Practices for Active Cluster

There are a couple of best practices worth following for active cluster installations:

- [Work with Large Amounts of Data, page 55](#)
- [Work with the Command Line Interface, page 55](#)

Work with Large Amounts of Data

When uploading large amounts of data into a cluster node, it is best to do it during off-hours. Some time-consuming operations might take longer than expected because of how updates are propagated. For example, importing a CAR file. A metadata change is applied first on the local node and then, if successful, propagated to all the other nodes in parallel.

Work with the Command Line Interface

When using the command line interface console for TDV Server and Active Cluster, make sure that text is not inadvertently selected or highlighted. This can cause cluster problems.

The cluster_util Command-Line Program

Active Cluster provides cluster_util to create and manage Active Clusters from the command line. This topic documents the subcommands available in the cluster_util program.

- [Using the cluster_util Program, page 57](#)
- [The -create Subcommand, page 60](#)
- [The -getConfig Subcommand, page 61](#)
- [The -join Subcommand, page 61](#)
- [The -remove Subcommand, page 62](#)
- [The -setClusterName Subcommand, page 62](#)

Using the cluster_util Program

The cluster_util command-line program lets you perform many of the actions you need for working with Active Cluster. You can:

- Create a cluster
- Get cluster configuration information
- Join a cluster
- Remove a server from a cluster
- Repair a cluster
- Set the cluster name

All cluster_util commands require administrator (MANAGE_SERVER) privileges.

To view the cluster_util command-line syntax

1. Run TDV Server on the computer for which you want configuration information.
2. Open a DOS window or UNIX terminal and navigate to the /bin directory of the TDV installation.
3. Type one of these commands:
cluster_util.bat (Windows)

Or
cluster_util.sh (UNIX)

The following table shows the syntax to perform each task:

To ...	Syntax
Create a cluster	<pre>cluster_util -server <hostname> [-port <port>] [-encrypt] -user <username> -password <password> [-domain <domain>] -create -clusterName <clustername> [-debug]</pre>
Get cluster configuration information	<pre>cluster_util -server <hostname> [-port <port>] [-encrypt] -user <username> -password <password> [-domain <domain>] -getConfig [-debug]</pre>
Join a cluster	<pre>cluster_util -server <target server hostname> [-port <target server port>] [-encrypt] -user <username> -password <password> [-domain <domain>] -join -memberServer <cluster server hostname> [-memberPort <cluster server port>] -memberUser <username> -memberPassword <password> [-memberDomain <domain>] [-debug]</pre>
Remove a server from a cluster	<pre>cluster_util -server <hostname> [-port <port>] [-encrypt] -user <username> -password <password> [-domain <domain>] -remove [-memberServer <target server hostname>] [-memberPort <target server port>] [-debug]</pre>
Set a cluster name	<pre>cluster_util -server <hostname> [-port <port>] [-encrypt] -user <username> -password <password> [-domain <domain>] -setClusterName -clusterName <clusterName> [-debug]</pre>
Repair a cluster	<pre>cluster_util -server <target server hostname> [-port <target server port>] [-encrypt] -user <username> -password <password> [-domain <domain>] -repair</pre>

The options for creating a new cluster are described in the following table.

Cluster Option	Description
-clustername <cluster_name>	The name of the new cluster. This name appears in the Manager on the CLUSTER MANAGEMENT page.
-create	Subcommand to create a new cluster.
-debug	Use this option to turn on debugging. In case of an exception, stack trace is enabled.
-domain <domain name>	The domain on which this server resides or the domain on which the server you are repairing resides. If creating a new cluster the domain on which the new cluster resides. Default is composite.
-getConfig	Subcommand to get cluster configuration information.
-join	Subcommand to join an existing cluster.
-memberDomain <domain name>	The domain on which the cluster server resides. Default is composite.
-memberPassword <password>	The password of the administrative user on the cluster server.
-memberPort <cluster server port>	The port of the cluster server to which you are joining. Default is 9400.
-memberPort <target server port>	The port of the TDV Server you are removing. Default=9400.
-memberServer <cluster server hostname>	Specify the hostname for a TDV Server that is already a member of the cluster you are joining. "localhost" is not a valid entry.
-memberServer <target server hostname>	Specify the hostname for a TDV Server that you are removing.
-memberUser <user name>	The name of an administrative user on the cluster server.
-password <password>	The password of the user creating a cluster. The password of the administrative user.

Cluster Option	Description
-password <password>	The password of the administrative user for this cluster.
-port <cluster server port>	The port of the server node to which you are connecting. Default is 9400.
-port <port>	The port defined for the TDV Server. Default is 9400.
-port <target server port>	The port of the server you are adding to the cluster. Default is 9400.
-remove	Subcommand to remove a TDV Server from a cluster.
-regroup	Subcommand to repair nodes in a cluster and set the timekeeper node.
-server <cluster server hostname>	Specify the hostname for a TDV Server in the cluster.
-server <cluster server hostname>	<p>Specify the hostname for a functional TDV Server node in the cluster.</p> <p>When using the remove command, this server is NOT the one you are removing with this command (unless it is the last server in the cluster, in which case you do not need to specify -memberServer or -memberPort), but it can be any other server that is a member of the cluster.</p> <p>When using the repair command, this server is NOT the one you are repairing with this command, but can be any other server that is a member of the cluster.</p>
-server <host_name>	Specify the hostname for the TDV Server that is to be the first server in the cluster.
-server <target server hostname>	Specify the hostname for a TDV Server that you are adding to the existing cluster.
-setClusterName	Command to change the cluster name.
-user <user name>	The name of an administrative user who has access to this TDV Server.

The -create Subcommand

You use the -create subcommand of cluster_util to create a new Active Cluster. Execute this command on the TDV Server that you want to be the first member of the cluster.

Example

```
./cluster_util.sh -server localhost -user admin -password admin
-create -clusterName DevCluster
```

This command creates a cluster with the name of DevCluster on a UNIX machine.

The -getConfig Subcommand

You use the -getConfig subcommand of cluster_util to display configuration information for an Active Cluster.

Example

```
./cluster_util.sh -server localhost -user admin -password admin
-getConfig
```

This command returns the configuration information for the cluster server node localhost.

The -join Subcommand

You use the -join subcommand of cluster_util to add a TDV Server to an existing Active Cluster.

Example

```
./cluster_util.sh -server localhost -port 9410 -user admin
-password admin
-join -memberServer localhost -memberPort 9400 -memberUser admin
-memberPassword admin
```

This command joins the TDV Server named localhost at port 9410 to the cluster to which the server named localhost at port 9400 belongs. Because a server can belong to only one cluster, no cluster name is required.

The -remove Subcommand

You use the -remove subcommand of cluster_util to remove a TDV Server from an Active Cluster. To remove an entire cluster, you remove all TDV Servers in the cluster, one at a time, using the same command. Removing the last server deletes the cluster.

Example

```
./cluster_util.sh -server localhost -port 9410 -user admin  
-password admin -remove
```

This command removes the TDV Server named localhost at port 9410 from the cluster to which it belongs. Because a server can belong to only one cluster, no cluster name is required.

The -setClusterName Subcommand

You use the -setClusterName subcommand of cluster_util to rename an Active Cluster.

Example

```
./cluster_util.sh -server localhost -user admin -password admin  
-setClusterName -clusterName ProdCluster1
```

This command renames the cluster with the name ProdCluster1.

Active Cluster SYS_CLUSTER System Table

The Studio resource tree displays the metadata stored as a set of system tables in a database called System. The system database tables are visible to all client applications that have the appropriate access privileges.

Studio users can only execute SQL SELECT statements on the system tables. They cannot change system table data (because changes could compromise performance).

The SYS_CLUSTER table provides the status of the Active Cluster to which this server belongs. Each cluster system has one row, including a row for the server to which you are connected, although the contact information is NULL for this server. You must have ACCESS_TOOLS and READ_ALL_STATUS rights to see the rows in the table.

The following table describes the schema for the SYS_CLUSTER table.

Column	TDV JDBC Data Type	Null able	Description
SERVER_ID	VARCHAR (255)		Server ID.
SERVER_NAME	VARCHAR (255)		Server name.
SERVER_HOST	VARCHAR (255)		Server hostname.
SERVER_PORT	INTEGER		Server port.
STATUS	VARCHAR (20)		Server status.
TIMEKEEPER	CHAR (1)		Y if the server is the timekeeper, otherwise N.
CONNECT_TIME	TIMESTAMP	Yes	The time the cluster connection was established by this server. Each time the server connects to a cluster node and becomes an active member of the cluster, this value is updated.
LAST_CONTACT	TIMESTAMP	Yes	The last time this server received a message across the cluster network connection.

Column	TDV JDBC Data Type	Nul lab le	Description
IDLE_DURATION	BIGINT		Number of milliseconds since LAST_CONTACT.
NUM_CONNECTIONS	BIGINT		Number of times this server successfully connected to the cluster.
BYTES_TO	BIGINT		Number of bytes sent to this server since it started.
BYTES_FROM	BIGINT		Number of bytes sent from the remote server to this server.
MESSAGE	VARCHAR (65535)	Yes	Additional information about the server.