

webMethods Integration Server Clustering Guide

Version 10.1

October 2017

This document applies to webMethods Integration Server Version 10.1 and to all subsequent releases.

Specifications contained herein are subject to change and these changes will be reported in subsequent release notes or new editions.

Copyright © 2007-2017 Software AG, Darmstadt, Germany and/or Software AG USA Inc., Reston, VA, USA, and/or its subsidiaries and/or its affiliates and/or their licensors.

The name Software AG and all Software AG product names are either trademarks or registered trademarks of Software AG and/or Software AG USA Inc. and/or its subsidiaries and/or its affiliates and/or their licensors. Other company and product names mentioned herein may be trademarks of their respective owners.

Detailed information on trademarks and patents owned by Software AG and/or its subsidiaries is located at <http://softwareag.com/licenses>.

Use of this software is subject to adherence to Software AG's licensing conditions and terms. These terms are part of the product documentation, located at <http://softwareag.com/licenses> and/or in the root installation directory of the licensed product(s).

This software may include portions of third-party products. For third-party copyright notices, license terms, additional rights or restrictions, please refer to "License Texts, Copyright Notices and Disclaimers of Third Party Products". For certain specific third-party license restrictions, please refer to section E of the Legal Notices available under "License Terms and Conditions for Use of Software AG Products / Copyright and Trademark Notices of Software AG Products". These documents are part of the product documentation, located at <http://softwareag.com/licenses> and/or in the root installation directory of the licensed product(s).

Table of Contents

About this Guide.....	5
Document Conventions.....	5
Online Information.....	6
An Overview of Clustering.....	7
What Is Clustering?.....	8
Load Balancing in a Cluster.....	9
What Data Is Shared in a Cluster?.....	9
Overview of Failover Support.....	9
Overview of Reliability.....	10
Guaranteed Delivery and Clustering.....	10
Checkpoint Restart.....	11
Putting It All Together.....	11
webMethods Integration Server Session Objects.....	13
Usage Notes for Session Storage.....	13
Scheduled Jobs.....	14
Client Applications.....	14
Remote Integration Servers in a Cluster.....	15
Installation and Setup.....	17
Overview.....	18
Considerations for Using Other webMethods Products in a Clustered Environment.....	18
Setting Up the webMethods Integration Servers in the Cluster.....	18
Server Environment.....	18
Switching from the Embedded Database to an External RDBMS.....	20
Using Terracotta as the Caching Software.....	21
Setting up the Terracotta Server Array for Use by the Cluster of Integration Servers.....	21
Enabling Clustering for an Integration Server.....	22
About Action on Startup Error Options.....	25
What Happens When Integration Server Cannot Connect to the Distributed Cache?.....	26
Controlling Logging for Caching.....	26
Scheduling Jobs to Run in the Cluster.....	27
Specifying Unique Logical Names for Integration Servers in a Cluster.....	27
Using Guaranteed Delivery with Clustering.....	28
Managing Server Clustering.....	31
Viewing the Servers in the Cluster.....	32
Removing a Server from the Cluster.....	32
Adding a Server Back into the Cluster.....	33
Taking a Server Offline.....	33
Bringing a Server Back Online.....	34

Increasing the Cache Size.....	35
Increasing the Maximum Elements on Disk Size.....	35
Increasing the Maximum Elements on Disk Size When the Terracotta Server Array Uses Permanent-Store Persistence Mode.....	35
Increasing the Maximum Elements on Disk Size When the Terracotta Server Array Uses Temporary-Swap-Only Persistence Mode.....	36
How Integration Server Cluster and Terracotta Server Array Handle Failures.....	36
What Happens when Integration Server Cannot Connect to the Terracotta Server Array?.....	37
Integration Server Cannot Download the tc-config.xml File.....	37
Solution.....	37
Integration Server Cannot Connect to the Servers in the tc-config.xml File.....	38
Solution.....	38
What Happens when Integration Server Is Disconnected from the Terracotta Server Array?.....	39
Solution.....	40
What Happens When Terracotta Servers in the Terracotta Server Array Are Disconnected?.....	40
How Do Integration Servers Respond when Terracotta Servers in the Terracotta Server Array Are Disconnected?.....	41
Solution.....	41
Index.....	43

About this Guide

This guide describes how to install and configure the webMethods Integration Server Clustering feature. It contains information for administrators who configure and manage a webMethods Integration Server and for application developers who want to create services that interact directly with the Integration Server short-term store.

Note: This guide describes features and functionality that may or may not be available with your licensed version of webMethods Integration Server. For information about the licensed components for your installation, see the **Settings > License** page in the Integration Server Administrator.

To use this guide effectively, you should understand the basic concepts described in the *webMethods Integration Server Administrator's Guide* and *webMethods Service Development Help*. You should also be familiar with all the servers you want to include in the cluster.

Document Conventions

Convention	Description
Bold	Identifies elements on a screen.
Narrowfont	Identifies storage locations for services on webMethods Integration Server, using the convention <i>folder.subfolder:service</i> .
UPPERCASE	Identifies keyboard keys. Keys you must press simultaneously are joined with a plus sign (+).
<i>Italic</i>	Identifies variables for which you must supply values specific to your own situation or environment. Identifies new terms the first time they occur in the text.
Monospace font	Identifies text you must type or messages displayed by the system.
{ }	Indicates a set of choices from which you must choose one. Type only the information inside the curly braces. Do not type the { } symbols.

Convention	Description
	Separates two mutually exclusive choices in a syntax line. Type one of these choices. Do not type the symbol.
[]	Indicates one or more options. Type only the information inside the square brackets. Do not type the [] symbols.
...	Indicates that you can type multiple options of the same type. Type only the information. Do not type the ellipsis (...).

Online Information

Software AG Documentation Website

You can find documentation on the Software AG Documentation website at <http://documentation.softwareag.com>. The site requires Empower credentials. If you do not have Empower credentials, you must use the TECHcommunity website.

Software AG Empower Product Support Website

You can find product information on the Software AG Empower Product Support website at <https://empower.softwareag.com>.

To submit feature/enhancement requests, get information about product availability, and download products, go to [Products](#).

To get information about fixes and to read early warnings, technical papers, and knowledge base articles, go to the [Knowledge Center](#).

Software AG TECHcommunity

You can find documentation and other technical information on the Software AG TECHcommunity website at <http://techcommunity.softwareag.com>. You can:

- Access product documentation, if you have TECHcommunity credentials. If you do not, you will need to register and specify "Documentation" as an area of interest.
- Access articles, code samples, demos, and tutorials.
- Use the online discussion forums, moderated by Software AG professionals, to ask questions, discuss best practices, and learn how other customers are using Software AG technology.
- Link to external websites that discuss open standards and web technology.

1 An Overview of Clustering

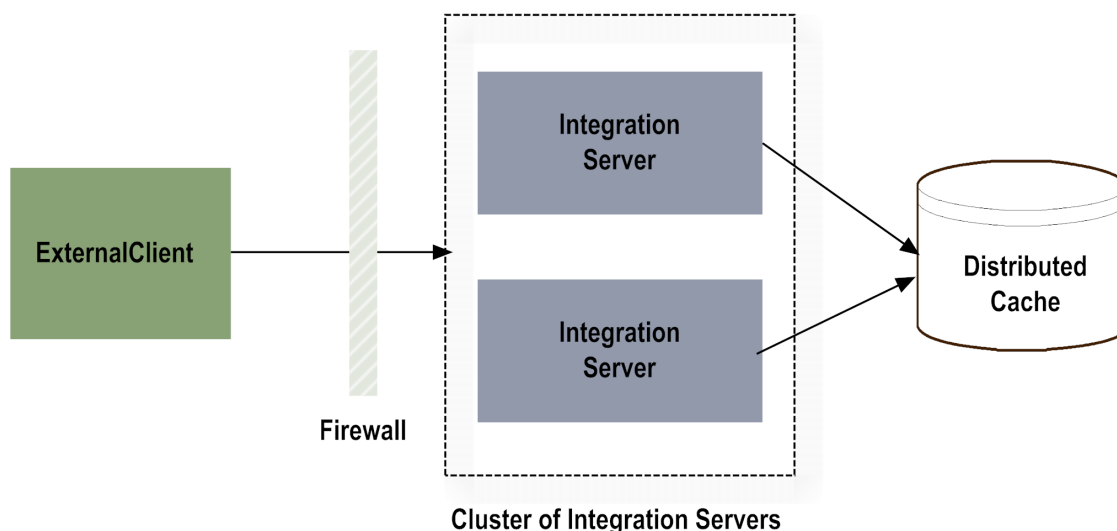
■ What Is Clustering?	8
■ What Data Is Shared in a Cluster?	9
■ Overview of Failover Support	9
■ Overview of Reliability	10
■ webMethods Integration Server Session Objects	13
■ Scheduled Jobs	14
■ Client Applications	14
■ Remote Integration Servers in a Cluster	15

What Is Clustering?

Clustering is an advanced feature of the webMethods product suite that substantially extends the reliability, availability, and scalability of the webMethods Integration Server. Clustering accomplishes this by providing the infrastructure and tools to deploy multiple Integration Servers as if they were a single virtual server and to deliver applications that leverage that architecture.

- **Scalability**—Without clustering, only vertical scalability is possible. That is, increased capacity requirements can only be met by deploying on larger, more powerful machines, typically housing multiple CPUs. Clustering provides horizontal scalability, which allows virtually limitless expansion of capacity by simply adding more machines of the same or similar capacity.
- **Availability**—Without clustering—even with expensive Fault-Tolerant systems—a failure of the system (hardware, Java runtime, or software) may result in unacceptable downtime. Clustering provides virtually uninterrupted availability by deploying applications on multiple Integration Servers; in the worst case, a server failure produces degraded but not disrupted service
- **Reliability**—Unlike a server farm (an independent set of servers), clustering provides the reliability required for mission-critical applications. Distributed applications must address network, hardware, and software errors that might produce duplicate (or failed) transactions. Clustering makes it possible to deliver "exactly once" execution as well as checkpoint/restart functionality for critical operations.

The following diagram shows clustering in its simplest form:



Important: Do not perform development work in a clustered environment. Basic namespace locking, the VCS Integration feature, and local service development are not supported across Integration Servers in a cluster.

Load Balancing in a Cluster

Load balancing is an optimizing feature you use with clustered webMethods Integration Servers. Load balancing controls how requests are distributed to the servers in the cluster. You must use a third-party load balancer to perform load balancing. A third-party load balancer can be useful if you need load balancing for multiple types of servers, for example, web servers and application servers, in addition to webMethods Integration Servers. Third-party load balancers also offer virtual IP support, but they are not "out of the box." Most third-party load balancers perform load balancing in a round-robin manner or based on network level metrics such as TCP connections and network response time.

What Data Is Shared in a Cluster?

In a clustered configuration, the session state for clients connected to all servers in the cluster is stored in the distributed cache created by the caching software. Integration Server uses Terracotta Server Array as the caching software. The cache allows transactions in a conversation to be continued on other servers in the cluster.

Information about scheduled jobs, jobs tracked by guaranteed delivery, and XREF data is stored in an external database, identified by the ISInternal functional alias. In addition, JOIN documents (ANDs and XORs only) and the processing state of activations that are participating in a JOIN are also stored in the database. For more information about functional aliases, refer to *Installing Software AG Products*.

Note: In a stateless cluster of Integration Servers, the session state of a client is not stored in a distributed cache; it is only available locally. This deployment model is appropriate for stateless services. For information about stateless clustering for Integration Servers, refer to *webMethods Integration Server Administrator's Guide*

Overview of Failover Support

Failover support allows you to recover from system failures that occur during processing, making your applications more robust. For example, by having more than one Integration Server, you protect your application from failure in case one of the servers becomes unavailable. If the webMethods Integration Server to which the client is connected fails, the client automatically reconnects to another webMethods Integration Server in the cluster.

Note: webMethods Integration Server clustering provides failover capabilities to clients that implement the webMethods Context and TContext classes. Integration Server does not provide failover capabilities when a generic HTTP client, such as a web browser, is used.

You can achieve failover support a number of ways:

- **webMethods Integration Server clustering.** With multiple webMethods Integration Servers, if one Integration Server fails, requests can be automatically redirected to another server, thereby avoiding a single point of failure.
- **Redundant hardware configurations.** By specifying hardware configurations that are redundant, for example multiple machines, multiple file servers, and so on, you further reduce the chance of a single point of failure.

Overview of Reliability

Several features of the webMethods Integration Server improve reliability in a clustered or unclustered environment.

- **Guaranteed delivery.** This feature ensures that a service executes once and only once. It is particularly useful when used with clustering to prevent a restarted service from running on more than one server. This feature is only for use with server-to-server communications.
- **Checkpoint restart services.** The `pub.storage` and `pub.cache` services allow you to code your flow service to store state information and other pertinent information in the short-term data store. If the flow service fails because a server becomes unavailable, the flow service can be restarted from the last checkpoint rather than at the beginning. For more information about `pub.storage` and `pub.cache` services, see the *webMethods Integration Server Built-In Services Reference*.

Guaranteed Delivery and Clustering

Guaranteed delivery ensures one-time execution of services by guaranteeing the following:

- Requests from the client to execute services are delivered to the server.
- Services are executed once, and only once.
- Responses from the execution of services are delivered to the client.

Guaranteed delivery is useful with or without clustering. If you are not using clustering, guaranteed delivery makes sure a client resubmits a service request until it succeeds and a response is returned. Guaranteed delivery also makes sure the service executes only once. For example, if the network connection between the client and the webMethods Integration Server fails after execution but before the response is successfully redirected to the client, the service might be executed twice. Guaranteed delivery prevents this from happening.

With clustering, if the server on which the service is running becomes unavailable, the client retries the service on another server in the cluster. If the request fails against all servers in the cluster, you must use guaranteed delivery to guarantee execution. As in an

unclustered environment, you must use guaranteed delivery to prevent a service from executing more than once.

For more information about guaranteed delivery, refer to *webMethods Integration Server Administrator's Guide* and the *Guaranteed Delivery Developer's Guide*.

Checkpoint Restart

webMethods Integration Server provides a number of built-in services (in the `pub.cache` and `pub.storage` folders) that you can use to make your flows more robust. With these services, you can write state information and other pertinent data to a data store in the Integration Server short-term data store. If the webMethods Integration Server on which your flow is executing becomes unavailable, when your flow restarts it can check the state information in the data store and begin processing at the point where the flow was interrupted. For more information about `pub.cache` and `pub.storage` services, see the *webMethods Integration Server Built-In Services Reference*.

Putting It All Together

This table summarizes how checkpoint restart, clustering, and guaranteed delivery work together to provide availability, failover, and reliability in different situations.

Checkpoint restart specified	Clustering specified	Guaranteed Delivery specified	If the server on which the service is running becomes unavailable	Point at which the service restarts
			After the server becomes available, you must manually restart the service.	At the beginning
		X	After the server becomes available, the client automatically restarts the service. The client keeps trying to run the service until it runs successfully.	At the beginning
	X		The client automatically retries the service on another server in the cluster. If that server fails, the client retries the service on the next server in the cluster, and so on. If all attempts fail, you	At the beginning

Checkpoint restart specified	Clustering specified	Guaranteed Delivery specified	If the server on which the service is running becomes unavailable	Point at which the service restarts
			must manually restart the service.	
	X	X	The client retries the service on the next server in the cluster. If that server fails, the client retries the service on the next server in the cluster, and so on. The client continues to try running the service until it runs successfully.	At the beginning
X			When the server becomes available again, you must manually restart the service.	Flow services: At the specified checkpoint Other services: At the beginning
X		X	When the server becomes available again, the client automatically retries the service. The client continues trying to run the service until it completes successfully.	Flow services: At the specified checkpoint Other services: At the beginning
X	X		The client automatically retries the service on the next server in the cluster. If that server fails, the client retries the service on the next server in the cluster, and so on. If attempts on all servers fail, you must restart the service manually.	Flow services: At the specified checkpoint Other services: At the beginning

Checkpoint restart specified	Clustering specified	Guaranteed Delivery specified	If the server on which the service is running becomes unavailable	Point at which the service restarts
X	X	X	The client automatically retries the service on the next server in the cluster. If that server fails, the client retries the service on the next server in the cluster, and so on. The client keeps trying until the service runs successfully.	Flow services: At the specified checkpoint Other services: At the beginning.

webMethods Integration Server Session Objects

Clustered Integration Servers create and maintain the session objects that are stored in a cache.

In a non-clustered environment, Integration Server maintains session information in its own local memory. In a cluster, however, Integration Server creates a session in the distributed (or shared) cache, so that when a load balancer redirects a client for failover, the new server can access the session information.

The Integration Server that initially receives a request from a client creates the session object in the cache. Other Integration Servers in the cluster can access the session object to access and update session information as necessary.

When you configure an Integration Server to use clustering, you specify a setting that indicates how long inactive session objects are maintained in the cache. Periodically, each Integration Server in the cluster checks the session objects in the cache to determine if any have expired, and if so, removes them.

Usage Notes for Session Storage

- Only objects that are serializable can be successfully stored in the session when Integration Server is running in a cluster. That is, those objects must implement the `java.io.Serializable` interface or one of the `com.wm.util.coder` interfaces such as `com.wm.util.coder.IDataCodable`. In a cluster, a session is serialized to the shared session store. When the session is restored from the shared store to an Integration Server, the complete state of the session, including any application data that had been saved to it, is available. If the session contains objects that are not serializable, those objects are converted into strings that hold the objects' class names. The actual state of those objects is lost.

The requirement that these objects be serializable applies to the entire object graph, including the object placed into the session and every object it contains, no matter how deeply nested.

Although your production application can run in a cluster, you will be developing it on a stand-alone Integration Server (clustered development is not supported). It is important to be aware of the serializable requirement so that you do not encounter problems with your session data once you start to test in a cluster.

- The processing speed of a cluster is determined in large part by network I/O. Adding application data to the session state will increase the amount of I/O the cluster must perform, and make it operate more slowly. The addition of a single large or complex object to each session can have a noticeable effect on the overall throughput of a cluster.

If you are concerned that saving application data to the session might impact the performance of your Integration Server cluster, consider other ways of saving this data. If it does not have to persist across server restarts and does not have to be shared throughout the cluster, a simple Java collection such as a Vector or HashMap might be appropriate. If the data needs to survive server restarts but does not have to be shared throughout the cluster, writing it to the local file system is an option. If the data needs to be shared throughout the cluster, consider saving it in a database.

- Even though sessions are created and maintained in a distributed cache, each Integration Server keeps a small portion of the cache locally. The local cache stores information relating to the sessions that are active on Integration Server as well as a list of nodes in the cluster. If you anticipate that your Integration Server sessions will use a large portion of the cache, you should increase the **Maximum Elements In Memory** or **Maximum Off-Heap** settings for each Integration Server in the cluster. For information about changing these settings, see *webMethods Integration Server Administrator's Guide*.

Scheduled Jobs

You can schedule jobs to run on one, any, or all Integration Servers in the cluster. Integration Server stores information about scheduled jobs in the external database identified by the ISInternal functional alias. For jobs to run in the cluster, the server must be enabled for clustering and existing jobs must be flagged to run in the cluster. For instructions, see the chapter about managing services in *webMethods Integration Server Administrator's Guide*.

Client Applications

Server clustering is almost transparent to the client. A client can issue requests to a server that is in a clustered environment in the same way it issues a request to a server that is not in a clustered environment.

When a client connects to a server that is in a clustered environment (using the `Context` class), the server returns information about the other servers in the cluster to the client. In the event that a request is not processed, the client can use this information to connect to another server in the cluster to have the request fulfilled.

Whether you are using the `Context` or `TContext` class to communicate with the webMethods Integration Server, the redirection of failed requests is transparent to the client. The logic to handle redirection is in the webMethods Integration Server `Context` or `TContext` class.

To improve the failover capability, before your client calls `Context` or `TContext` to connect to an Integration Server in the cluster, have your client issue the `setRetryServer` method in that class to specify another server to try if the first server the client tries to connect to is unavailable.

No special processing is required in your clients.

Note: webMethods Integration Server clustering provides failover capabilities to HTTP-based webMethods clients, such as those clients built using the webMethods `Context` and `TContext` classes.

Remote Integration Servers in a Cluster

An Integration Server can be configured to connect to a remote server for a number of reasons, including:

- Allow clients to run services on other Integration Servers using the `pub.remote:invoke` service and the `pub.remote.gd:*` services.
- Connect publisher and subscriber Integration Servers to each other for the purpose of package replication.
- Facilitate the process of presenting different certificates to different Integration Servers.

In general, if a remote server is not available when a request is made, and the remote server is not part of a cluster, Integration Server will use the retry server specified in the alias definition for the remote server.

If the requested remote server is part of a cluster, Integration Server will attempt to use the next Integration Server in the cluster until one is found. If no available Integration Server is found in the cluster, Integration Server will use the retry server specified in the alias definition for the remote server.

In cases where a client calls the `pub.remote:invoke` to run a service on a remote server that is part of a cluster, it is possible to modify the default behavior by using the `$clusterRetry`.

This parameter controls whether the service will try other Integration Servers in the cluster. When this parameter is set to true, the service will try other Integration Servers in the cluster. If none are found, the service will try the retry server specified in the remote server alias definition. When this parameter is set to false, the service does not

try other Integration Servers in the cluster. Instead, it will try the retry server specified in the alias definition.

For more information about remote servers, see *webMethods Integration Server Administrator's Guide*.

For more information about the `pub.remote:invoke` service, see the *webMethods Integration Server Built-In Services Reference*.

You can also use the Java API to control the retry behavior. Specifically, if you are using the `Context` or `TContext` class, you can control whether Integration Server looks for a retry server by calling the `BaseContext.setAllowRedir` method. In addition, you can specify which retry server to use by calling the `BaseContext.setRetryServer` method. For more information about the `Context` and `TContext` classes, see the *webMethods Integration Server Java API Reference*.

2 Installation and Setup

■ Overview	18
■ Setting Up the webMethods Integration Servers in the Cluster	18
■ Using Guaranteed Delivery with Clustering	28

Overview

This section summarizes the steps you must take to set up server clustering:

- **Install the Integration Servers.** All Integration Servers that you want to include in the cluster should be running the same version of Integration Server software and be at the same patch level.
- **Install and configure the Terracotta Server Array.** You must install and configure the Terracotta Server Array *before* you enable clustering on the Integration Servers. For more information about this step, see ["Setting up the Terracotta Server Array for Use by the Cluster of Integration Servers"](#) on page 21.
- **Configure the Integration Server for clustering.** Instructions are provided in ["Setting Up the webMethods Integration Servers in the Cluster"](#) on page 18.

Considerations for Using Other webMethods Products in a Clustered Environment

Other webMethods products support the use of clustering to synchronize requests, transactions, notifications, and so on across Integration Server instances in the cluster. For more information about whether a particular product supports clustering and how to configure that product for a clustered environment, see the documentation provided for the product.

Setting Up the webMethods Integration Servers in the Cluster

It is recommended that you maintain the same environment for all servers in a cluster, for example, the same packages of services and ACLs that protect services. For information on the server environment variables that you should maintain, refer to ["Server Environment"](#) on page 18 below.

You can have a variety of Integration Servers in your cluster, for example one server on Windows, one on UNIX, and one on Linux.

Server Environment

Although not required, it is recommended that the servers in a cluster each have the same server environment. For server clustering to work effectively, you should keep the following the same on all servers in the cluster:

- **Each server should have the same set of licensed capabilities.**
- **All servers should have the same packages of services.** For a service to execute on any server in the cluster, that service must exist in the same package on every server in the cluster.

Note: Software AG recommends that you use webMethods Deployer, or the package replication functionality in the Integration Server Administrator to copy packages to other servers in the cluster, instead of using Designer to copy them. For information about webMethods Deployer, see the *webMethods Deployer User's Guide*. For information about package replication, see *webMethods Integration Server Administrator's Guide*.

- **Each server should have the same user accounts.** The server uses user account information to authenticate clients. When a server redirects a request to an alternate server, the alternate server re-authenticates the user using the credentials supplied to the original server. If authentication fails, the request fails.
- **Access to services should be the same on all servers.** Integration Server uses group information and ACLs to determine whether a client has access to a service. All Integration Servers should have the same groups with the same group membership. Services should be protected by the same ACLs. The ACLs should identify the same Allow Groups and Deny Groups.

If a request is redirected to a server that denies access to the requested service, the request will fail.

- **Each server should have the same public caches.** One of the purposes of clustering is to allow a session to be directed to any server within a cluster and for the session to be processed the same way. Some of that processing may involve caching, in which case the same caches, with identical configurations, need to be on each server in the cluster.

Note: Software AG recommends that you use webMethods Deployer to copy cache managers and caches to other servers in the cluster. For information about webMethods Deployer, see the *webMethods Deployer User's Guide*.

- **Each server should have the same event subscriptions.** Integration Server saves information for event types and event subscriptions in the eventcfg.bin file. This file is generated the first time you start an Integration Server and is located in the following directory: *Integration Server_directory/instances/instance_name/config*. Copy this file from one server to another to duplicate event subscriptions on all servers in the cluster.
- **Clock time must be the same on all servers.** The clocks on all machines in the Integration Server cluster must be synchronized for scheduled jobs to work properly in a cluster.
- **Each Integration Server should connect to the same messaging providers.** All servers should use the same messaging connection aliases to connect to the same messaging providers (Broker and/or Universal Messaging). Each messaging connection alias with the same name must use the same client prefix. Furthermore, each Broker connection alias must use the same client group.
- **Each server should connect to the same set of databases.** For example, if you are using the audit database to store audit data, all servers must connect to the same audit database. If you are using a back-end database for application-related data, all servers must connect to the same back-end database.

- **Each server should have its own diagnostic port.** If you plan to troubleshoot using the diagnostic port, you must configure a diagnostic port for each server in the cluster. The diagnostic port can access only the Integration Server on which it is defined. For more information about the diagnostic port, see *webMethods Integration Server Administrator's Guide*.
- **Each server should have its own Tspace.** If you want the Integration Servers in a cluster to temporarily store large documents in a hard disk drive space rather than keep them in memory, you must define a different Tspace for each Integration Server in a cluster.
- **Each server must have the same cluster name.**
- **Every server must point to the same Terracotta Server Array.** The Integration Servers must use the same Terracotta Server Array URLs and the same cluster name. Additionally, the Terracotta Server Array that you want to use to store the system caches must already be configured. Session timeout and time-to-live should be the same on all servers. If it is not the same, session lifetime will be unpredictable. For more information about using Terracotta with the cluster, see "[Using Terracotta as the Caching Software](#)" on page 21.

Switching from the Embedded Database to an External RDBMS

For Integration Servers to participate in a cluster, they must share database components in an external RDBMS. If you installed an Integration Server with the embedded database, but now want to add it to a cluster, you must switch the Integration Server to use the shared external RDBMS.

To switch to a shared external RDBMS

1. In Integration Server Administrator, navigate to the **Security > Certificates > Client Certificates** screen, and make a note of the certificate mappings.

If the cluster already exists, create JDBC connection pools for the IS Internal and Cross Reference database components, and point the ISInternal and Xref functions at the shared IS Internal and Cross Reference database components. See *Installing Software AG Products* for instructions.

If those database components do not yet exist, create them and connect all Integration Servers in the cluster to them.

2. Run the migration utility `pub.scheduler:migrateTasksToJDBC` to migrate your scheduled tasks from the embedded database to the external RDBMS. See the *webMethods Integration Server Built-In Services Reference* for instructions.

Note: This service migrates scheduled tasks only; client mappings and run-time data will not be migrated.


3. In Integration Server Administrator, navigate to the **Security > Certificates > Client Certificates** screen and re-specify your certificate mappings. For instructions

on mapping a client certificate to a user, refer to *webMethods Integration Server Administrator's Guide*.

4. Configure Integration Server as described below.

Using Terracotta as the Caching Software

The data that the Integration Servers share is stored in distributed caches on a Terracotta Server Array. A Terracotta Server Array is a group of one or more Terracotta Servers. The data associated with a cache is spread across the servers in the array (each server maintains a portion of the cache). You can mirror the servers in the array for high availability. For more information about Terracotta Server Arrays, see the Terracotta product documentation.



You must install and configure the Terracotta Server Array *before* you enable clustering on the Integration Servers. For the list of general steps that you must perform before you enable clustering, see ["Setting up the Terracotta Server Array for Use by the Cluster of Integration Servers"](#) on page 21.

Setting up the Terracotta Server Array for Use by the Cluster of Integration Servers

Before you enable clustering using Terracotta, perform the following steps to ensure that the Terracotta Server Array is ready for use by the Integration Servers in the cluster.

Step	Description
1	Install the Terracotta Server Array if you have not done so already. For installation instructions, see <i>Installing Software AG Products</i> .

Note: Before you install your Terracotta Server Array, you need to decide how many Terracotta Servers will make up the array and whether or not you will mirror those servers. To guide your decision, review the sections on high availability and architecture in the product

Step	Description
	documentation for the Terracotta Server Array or consult your Software AG representative.
2	Configure the tc-config.xml file on the Terracotta Server Array as described in the section on configuring the tc-config.xml file in the Ehcache chapter of <i>webMethods Integration Server Administrator's Guide</i> . This section describes the parameter settings that are required when using a Terracotta Server Array with Integration Server.
3	Install your Terracotta license key on each Integration Server in the cluster. For procedures, see the section on installing and changing a Terracotta license in the Ehcache chapter of <i>webMethods Integration Server Administrator's Guide</i> .

After you complete these steps, you can enable clustering as described in "[Enabling Clustering for an Integration Server](#)" on page 22.

Note: The minimum offheap (or memory) required for clustering in webMethods hybrid mode is 2GB. However, for non-hybrid deployments, distributed caching and other applications, refer to the Terracotta documentation for recommended configuration

Enabling Clustering for an Integration Server

After you ensure that Integration Server has the appropriate environment, you can add it to the cluster by configuring the server to use clustering.

Keep the following points in mind when enabling clustering for an Integration Server.

- You must complete the setup steps described in "[Setting up the Terracotta Server Array for Use by the Cluster of Integration Servers](#)" on page 21 before you enable clustering on the Integration Servers.
- To be in the same cluster, Integration Servers must use the same Terracotta Server Array URLs and the same cluster name.
- An enterprise can have more than one cluster. To isolate multiple clusters on the same network, each cluster must have a different cluster name or different Terracotta Server Array or both.
- You must have webMethods Integration Server administrator privileges to enable clustering. If you do not have administrator privileges, have your webMethods Integration Server administrator perform this procedure.

To enable clustering

1. In the **Settings** menu of the navigation area, click **Clustering**.

2. Click **Edit Cluster Settings**.
3. Click **Enable Cluster**.
4. Specify the following information:

For this parameter	Specify				
Cluster Name	<p>Name of the cluster to which this Integration Server belongs.</p> <p>Keep the following in mind when specifying the cluster name:</p> <ul style="list-style-type: none"> ■ The cluster name cannot include any periods “.”. Integration Server converts any periods in the name to underscores when you save the cluster configuration. ■ The cluster name cannot exceed 32 characters. If it does, Integration Server uses the first 20 characters of the supplied name and then a hash for the remaining characters. 				
Session Timeout	<p>Number of minutes an inactive session will be retained in the clustered session store. The default is 60.</p> <p>Set the clustered session timeout value to be longer than the session timeout value, which governs how long a server keeps session information in its memory. (Integration Server Administrator displays the session timeout on the Settings > Resources screen.)</p>				
Action on Startup Error	<p>How Integration Server responds when an error at start up prevents Integration Server from joining the cluster. Select one of the following</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-decoration: underline;">Option</th> <th style="text-decoration: underline;">Description</th> </tr> </thead> <tbody> <tr> <td style="vertical-align: top;">Start as Stand-Alone Integration Server</td> <td>Integration Server starts as a stand-alone, unclustered Integration Server if it encounters any errors that prevent it from joining the cluster at start up. Integration Server continues to receive requests on inbound ports and serve requests. This is the default.</td> </tr> </tbody> </table>	Option	Description	Start as Stand-Alone Integration Server	Integration Server starts as a stand-alone, unclustered Integration Server if it encounters any errors that prevent it from joining the cluster at start up. Integration Server continues to receive requests on inbound ports and serve requests. This is the default.
Option	Description				
Start as Stand-Alone Integration Server	Integration Server starts as a stand-alone, unclustered Integration Server if it encounters any errors that prevent it from joining the cluster at start up. Integration Server continues to receive requests on inbound ports and serve requests. This is the default.				

For this parameter	Specify
Shut Down Integration Server	Integration Server shuts down if it encounters any errors that prevent it from joining the cluster at start up.
Enter Quiesce Mode on Stand-Alone Integration Server	Integration Server starts as a stand-alone, unclustered Integration Server in quiesce mode if it encounters any errors that prevent it from joining the cluster at start up. When Integration Server is in quiesce mode, only the diagnostic port and quiesce port are enabled.
	For more information about the Action on Startup Error options, see " About Action on Startup Error Options " on page 25 .
Terracotta Server Array URLs	A comma separated list of the URLs for the Terracotta Server Array to be used with the cluster to which this Integration Server belongs. The URLs must be in the following format: <i>host :port</i>

5. Click **Save Settings**.
6. Restart the Integration Server.
7. Using Integration Server Administrator, navigate to **Settings > Cluster** and verify that all nodes in the cluster are displayed under **Cluster Hosts**.

Notes:

- If you specify a name that does not already exist, Terracotta Server Array creates a new cache manager for each system cache manager on Integration Server. Terracotta Server Array appends the cluster name to the name of the system cache manager. System cache managers begin with `SoftwareAG`. For example, if you name the cluster "myCluster" the system cache manager `SoftwareAG.IS.Core` becomes `SoftwareAG.IS.Core.myCluster`.
- If you specify the name of a cluster that already exists and specify the same Terracotta Server Array URL used by the existing cluster name, the caching software adds this Integration Server to that cluster.

About Action on Startup Error Options

When Integration Server that is configured for clustering starts up, errors may prevent Integration Server from connecting to the Terracotta Server Array. These errors can be caused by some of the following: a configuration issue with the Terracotta license file, an unavailable or incorrectly configured Terracotta Server Array, or an unavailable database.

When you configure clustering for an Integration Server, you can determine how Integration Server responds when the server starts and cannot connect to the Terracotta Server Array. The action Integration Server takes affects whether Integration Server is available to process requests, which Integration Server features are available, and how you might correct the errors. The **Action on Startup Error** parameter determines how Integration Server reacts when the Terracotta Server Array cannot be reached at startup. You can specify one of the following options for the **Action on Startup Error** parameter.

- **Start as Stand-Alone Integration Server.** If Integration Server encounters any errors that prevent it from joining the cluster at start up, Integration Server starts as a stand-alone, unclustered Integration Server. Integration Server continues to receive requests on inbound ports and serve requests. Integration Server does not use any distributed caches and instead uses local caches. You can use Integration Server to make changes needed to resolve the startup errors. When you restart Integration Server, Integration Server, will rejoin the cluster if Integration Server does not encounter any start up errors that prevent the connection to the Terracotta Server Array.

Start as Stand-Alone Integration Server is the default option for the **Action on Startup Error** parameter.

- **Shut Down Integration Server.** If Integration Server encounters any errors that prevent it from joining the cluster at start up, Integration Server shuts down immediately. To make any changes to the clustering configuration, you must edit the configuration while running in safe mode. If you need to change the Terracotta Server Array URLs, you must start Integration Server in safe mode, change the values, and restart the Integration Server.

For more information about starting Integration Server in safe mode, see *webMethods Integration Server Administrator's Guide*

- **Enter Quiesce Mode on Stand-Alone Integration Server.** If Integration Server encounters any errors that prevent it from joining the cluster at start up, Integration Server starts as a stand-alone, unclustered Integration Server in quiesce mode. When Integration Server is in quiesce mode, only the diagnostic port and quiesce port are enabled. Integration Server will not receive requests on inbound ports. However, you can use Integration Server and Integration Server Administrator to make changes needed to correct the startup errors. After correcting the errors, exit quiesce mode. When exiting quiesce mode, Integration Server starts all the cache managers and joins the cluster.

If Integration Server encounters any errors that prevent it from joining the cluster when exiting quiesce mode, Integration Server does not re-enter quiesce mode. Instead, Integration Server runs as an unclustered, stand alone server. Integration Server captures any errors that occurred while attempting to join the cluster in a report displayed in Integration Server Administrator when the server exits quiesce mode. Use the report to help correct the underlying configuration errors. Then, to rejoin the cluster, either restart Integration Server or enter and then exit quiesce mode.

To use the **Enter Quiesce Mode on Stand-Alone Integration Server** option, a quiesce port must be configured for Integration Server. If a quiesce port is not configured and the **Enter Quiesce Mode on Stand-Alone Integration Server** option is selected, when Integration Server encounters any errors that prevent it from joining the cluster at start up, Integration Server shuts down instead of entering quiesce mode.

For more information about quiesce mode, see *webMethods Integration Server Administrator's Guide*

What Happens When Integration Server Cannot Connect to the Distributed Cache?

When Integration Server cannot connect to the distributed cache, one of the following occurs:

- If Integration Server cannot connect to the distributed cache at the time the Integration Server initializes, Integration Server logs an error to the server log and takes the action associated with the selected option for the **Action on Startup Error** parameter. To enable clustering, diagnose and correct the problem. For more information about the **Action on Startup Error** parameter, see "[About Action on Startup Error Options](#)" on page 25.
- If Integration Server becomes disconnected from the distributed cache after initialization completes, the rejoin behavior instructs the system cache managers on the Integration Server to reconnect to the Terracotta Server Array automatically. All system cache managers are configured to rejoin. For more information about the rejoin behavior for distributed cache managers, see *webMethods Integration Server Administrator's Guide*.

Controlling Logging for Caching

Logging for cache activity on the Terracotta Server Array is controlled by Ehcache. Ehcache logging has a default configuration, but you can change it by modifying the `.tc.dev.log4j.properties` file. This file is located in the *Integration Server_directory*. Ehcache creates additional log files on Integration Server when Integration Server connects to a Terracotta Server Array. You can specify where you want Ehcache to put the log files by setting the `watt.server.cachemanager.logsDirectory` property on Integration Server. The default value for this property is *Integration Server_directory/instances/instance_name/logs/tc-client-logs* directory. For information about how

Ehcache logs cache manager activity when Integration Server connects to a Terracotta Server Array, see *webMethods Integration Server Administrator's Guide*.

Scheduling Jobs to Run in the Cluster

You can schedule jobs to run on one, any, or all Integration Servers in the cluster. For jobs to run in the cluster, the server must be enabled for clustering and existing jobs must be flagged to run in the cluster. For instructions, see the chapter about scheduling services in *webMethods Integration Server Administrator's Guide*.

Specifying Unique Logical Names for Integration Servers in a Cluster

By default, Integration Server uses the host name to identify itself while scheduling tasks. However, when a cluster of Integration Servers are hosted on a single machine, the host name itself cannot uniquely identify the individual Integration Server instances. In such cases, use the `watt.server.scheduler.logical.hostname` property to specify a unique logical name to identify individual Integration Server instances.

The default value for this parameter is the host name.

Keep the following points in mind when setting the `watt.server.scheduler.logical.hostname` property:

- Set this property only if you are running multiple Integration Server in a cluster on a single machine.
- Set this property on each Integration Server instance in the cluster before scheduling any tasks.
- The logical host name you specify must be unique.
- The logical host name you specify cannot contain a semicolon (;).

To specify a unique name for an Integration Server in a cluster

1. Open the Integration Server Administrator if it is not already open.
2. In the **Settings** menu of the Navigation panel, click **Extended**.
Integration Server displays a screen that lists the configuration properties.
3. Locate the `watt.server.scheduler.logical.hostname` parameter. If this parameter does not exist, add it. Set it as follows:

Set this parameter...

```
watt.server.scheduler.  
logical.hostname
```

To specify...

A unique logical name for Integration Server.

For more information about using the Extended Settings screen to configure the Integration Server, see *webMethods Integration Server Administrator's Guide*.

4. Click **Save Changes**.
5. Restart Integration Server for the changes to take effect.

Using Guaranteed Delivery with Clustering

The guaranteed delivery capabilities of the webMethods Integration Server ensure guaranteed one-time execution of services. Guaranteed delivery ensures the following:

- Requests from the client to execute services are delivered to the server.
- Services are executed once, and only once.
- Responses from the execution of services are delivered to the client.

In a clustered environment, if the Integration Server on which the service is running becomes unavailable, the client retries the service on another Integration Server in the cluster. However, to make sure a request will run if it fails on every server in the cluster, you must configure your Integration Servers to use guaranteed delivery. In addition, as in an unclustered environment, you must use guaranteed delivery to prevent a service from executing more than once. For more information about guaranteed delivery, refer to *webMethods Integration Server Administrator's Guide*. For more information about developing services that use guaranteed delivery, refer to the *Guaranteed Delivery Developer's Guide*.

When you use webMethods Integration Server clustering, guaranteed delivery stores information about requests for *all* the servers in the cluster in a shared, centrally located database. Because the information is stored centrally, guaranteed delivery uses a locking mechanism to synchronize updates to the database.

Attempts to update the database might time out. As a result, the server may require several attempts to complete a request. You can limit how long a cluster server sleeps between attempts to place an update lock on the database. In addition, you can set a maximum time limit after which the server overrides a lock on the database held by a non-responsive server in order to complete the update.

To configure cluster database locking

1. Open the Integration Server Administrator if it is not already open.
2. In the **Settings** menu of the Navigation panel, click **Extended**.
The server displays a screen that lists configuration properties specified in the server.cnf file.
3. To set how long a cluster server waits between attempts to place an update lock, locate the **watt.server.tx.cluster.lockTimeoutMillis** parameter. If this parameter does not exist, add it. Set it as follows:

Set this parameter:

```
watt.server.tx.cluster.  
lockTimeoutMillis
```

To specify

The number of milliseconds a cluster server waits between requests to place an update lock.

4. To set how long a cluster server will wait before overriding a lock, locate the **watt.server.tx.cluster.lockBreakSecs** parameter. If this parameter does not exist, add it. Set it as follows:

Set this parameter:

```
watt.server.tx.cluster.  
lockBreakSecs
```

To specify

The number of seconds a cluster server waits before overriding a lock.

For more information about using the Extended Settings screen to configure the Integration Server, see *webMethods Integration Server Administrator's Guide*.

5. Click **Save Changes**.
6. Restart the server to put your changes into effect.

3 Managing Server Clustering

■ Viewing the Servers in the Cluster	32
■ Removing a Server from the Cluster	32
■ Adding a Server Back into the Cluster	33
■ Taking a Server Offline	33
■ Bringing a Server Back Online	34
■ Increasing the Cache Size	35
■ How Integration Server Cluster and Terracotta Server Array Handle Failures	36

Viewing the Servers in the Cluster

When you have server clustering enabled, you can display a list of all the servers in the cluster. The server learns of other servers in the cluster by retrieving information from the distributed (or shared) cache.

To view a list of all clustered servers

1. Start the Integration Server Administrator. See *webMethods Integration Server Administrator's Guide* if you need help with this step.
2. In the **Settings** menu of the navigation area, click **Cluster**.

The server displays the list of servers in the cluster.

If you notice that a server is missing from the list, it might be for one of the following reasons:

- The server is not running.
- The server cannot connect to the Terracotta Server Array.
- The clock time on the servers does not match.

Removing a Server from the Cluster

If you no longer want a server to be part of a cluster, disable clustering on that server.

To remove a server from the cluster

1. If you are using guaranteed delivery, you must take the server offline before removing it from the cluster. If you do not, the server will continue to receive transactions, but the transactions will not be integrated into the database and may be lost if the server is restored to the cluster. See "[Taking a Server Offline](#)" on page 33 for instructions on taking a server offline.
2. If it is not already running, start the Integration Server Administrator. See *webMethods Integration Server Administrator's Guide* if you need help with this step.
3. In the **Settings** menu of the navigation area, click **Clustering**.
4. Click **Edit Cluster Settings**.
5. Click **Disable Cluster** to turn clustering off for this server.
6. Click **Save Settings**.
7. Restart the server.

Note: Until you restart the server, it will remain in the list of cluster hosts on other Integration Servers in the cluster.

Adding a Server Back into the Cluster

You can add a server back into a cluster by enabling clustering on the server again. When clustering was previously enabled, the server stored the configured settings in the `server.cnf` file in the `Integration Server_directory/instances/instance_name/config` directory.

Configuration information is stored on the Terracotta Server Array. If you re-enable clustering and specify the same **Terracotta Server Array URLs** and the same cluster name, Integration Server retrieves the cache configuration information from the Terracotta Server Array.

When you enable clustering, ensure that the server has the same licensed capabilities as the rest of the servers in the cluster.

To add a server back into the cluster

1. If it is not already running, start the Integration Server Administrator of the server you want to add back into the cluster. See *webMethods Integration Server Administrator's Guide* if you need help with this step.
2. In the **Settings** menu of the navigation area, click **Clustering**.
3. Click **Edit Cluster Settings**.
4. Click **Enable Cluster** to turn clustering on for this server.
5. Make any changes you want to the configuration.
6. Click **Save Settings**.
7. Restart the server.

Note: Until you restart the server, it will not appear in the list of cluster hosts on other Integration Servers in the cluster.

Important: If you are using guaranteed delivery, the server you are returning to the cluster might be offline. When a server is offline, it can only be accessed through a single, unpublished port. Before you can use the server, you must first bring it back online. See "[Bringing a Server Back Online](#)" on page 34 for instructions.

Taking a Server Offline

There may be times when you need to take a server offline. When you take a server offline, you are limiting access to it through a single, unpublished port.

Before taking a server offline, make sure that the Integration Server has an existing, unpublished port. If one does not exist, create it.

Note: If you have guaranteed delivery, you must take a server offline before removing it from the cluster. If you do not, the server will continue to receive transactions, but the transactions will not be integrated into the database and may be lost if the server is restored to the cluster.

To take a server offline

1. Start the Integration Server Administrator. See *webMethods Integration Server Administrator's Guide* if you need help with this step.
2. In the **Security** menu of the navigation area, click **Ports**.
3. Click **Change Primary Port**.
4. In the **Select New Primary Port** area of the screen, select the port that you want to make the primary port. Typically, this is an existing, unpublished port.
5. Click **Update**.
6. On your browser, update the URL for the Integration Server Administrator to use the new port.
7. Disable all other listening ports.

Bringing a Server Back Online

There may be times when you need to bring a server back online. For example, you might have taken a server offline because you run guaranteed delivery and you needed to remove the server from the cluster. (When a server is offline, it can only be accessed through a single, unpublished port.) After you add the server back to the cluster, you must bring it back online.

To bring a server back online

1. Start the Integration Server Administrator using the new listening port that you designated as the primary port when you took the server offline. See *webMethods Integration Server Administrator's Guide* if you need help with this step.
2. In the **Security** menu of the navigation area, click **Ports**.
3. Re-enable the port that you disabled when you took the server offline.
4. Click **Change Primary Port**.
5. Change the primary listening port back to the published primary port for the server.
6. Click **Update**.
7. On your browser, update the URL for the Integration Server Administrator to the published primary port. This is the original primary port used by the server before you took the server offline.
8. Enable all disabled listening ports.

Increasing the Cache Size

If you anticipate that your Integration Server sessions will use a large portion of the cache, you should increase the cache size on each server in the cluster. You can increase the on-heap cache size by modifying the **Maximum Elements on Disk** setting. Or, if you are using BigMemory, you can modify the **Maximum Off-Heap** setting. For information about how to increase the **Maximum Elements On Disk** size, see "[Increasing the Maximum Elements on Disk Size](#)" on page 35. For information about increasing the **Maximum Off-Heap** size, see the Ehcache chapter of the *webMethods Integration Server Administrator's Guide*.

Increasing the Maximum Elements on Disk Size

In a clustered environment, the **Maximum Elements on Disk** cache setting is non-dynamic. The procedure you use to modify the value depends on whether your Terracotta Server Array is configured to use a persistence mode of permanent-store or temporary-swap-only (the default). The following sections provide instructions about how to increase the **Maximum Elements on Disk** size based on the persistence mode. For information about how to determine the persistence mode of the Terracotta Server Array, see the Ehcache product documentation for 2.8 at <http://ehcache.org/documentation>.

Increasing the Maximum Elements on Disk Size When the Terracotta Server Array Uses Permanent-Store Persistence Mode

To increase the Maximum Elements on Disk size when the Terracotta Server Array uses permanent-store persistence mode

1. Shut down all Integration Servers connected to the Terracotta Server Array.
2. On each Integration Server, open the SoftwareAG-IS-Core.xml file in a text editor. You can find the file in the following location:
Integration Server_directory/instances/instance_name /config/Caching
3. Increase the value of the `maxEntriesLocalDisk` parameter for the cache you want to change. The default value for each cache is 1000.

Note: The value for `maxEntriesLocalDisk` must be the same for each Integration Server in the cluster.

4. Save the files.
5. Shut down all servers in the Terracotta Server Array.
6. On each Terracotta Server Array server, open the tc-config.xml file in a text editor. You can find the tc-config.xml file in the following location:

TerracottaHome \bin

7. Locate the `<servers><server><data>` element in the `tc-config.xml` file and delete it.

Important: The `<servers><server><data>` element indicates the directory where the cache configuration and cached data reside.

8. Save the files.
9. Start all the servers in the Terracotta Server Array.
10. Start one Integration Server with the new configuration and verify that the change has taken effect. You can do so from the **Settings > Caching** pages in Integration Server Administrator.
11. Start the other Integration Servers.

Increasing the Maximum Elements on Disk Size When the Terracotta Server Array Uses Temporary-Swap-Only Persistence Mode

To increase the Maximum Elements on Disk size when the Terracotta Server Array uses temporary-swap-only persistence mode

1. Shut down all Integration Servers connected to the Terracotta Server Array.
2. On each Integration Server, open the `SoftwareAG-IS-Core.xml` file in a text editor. You can find the file in the following location:
Integration Server_directory/instances/instance_name /config/Caching
3. Increase the value of the `maxEntriesLocalDisk` parameter for the cache you want to change. The default value for each cache is 1000.

Important: The value for `maxEntriesLocalDisk` must be the same for each Integration Server in the cluster.

4. Save the files.
5. Shut down and restart all servers in the Terracotta Server Array.
6. Start one Integration Server with the new configuration and verify that the change has taken effect. You can do so from the **Settings > Caching** pages in Integration Server Administrator.
7. Start the other Integration Servers.

How Integration Server Cluster and Terracotta Server Array Handle Failures

This section contains information for the server administrator who configures and troubleshoots the Integration Server cluster. Consider the failure modes and information described in this section when creating your start-up scripts and procedures. Your

scripts and procedures should take into account how Integration Server responds if a failure occurs.

In Terracotta, you can set failover behavior to consistency instead of availability (the default). For information, see the section on failover tuning for guaranteed consistency in the *BigMemory Max Administrator Guide*.

What Happens when Integration Server Cannot Connect to the Terracotta Server Array?

How Integration Server behaves when it cannot connect to the Terracotta Server Array depends on whether Integration Server has made the initial connection to the Terracotta Server Array and has obtained the `tc-config.xml` file. When Integration Server starts up, it connects to the first specified Terracotta Server Array server in the **Terracotta Server Array URLs** list and downloads the `tc-config.xml` file. Integration Server then uses the settings in the `tc-config.xml` file to make connections to the servers in the Terracotta Server Array.

For more information about the `tc-config.xml` file, see *Using Terracotta with webMethods Products* and *webMethods Integration Server Administrator's Guide*.

Integration Server Cannot Download the `tc-config.xml` File

If **Terracotta Server Array** Integration Server is unable to download the `tc-config.xml` file, the startup sequence pauses while Integration Server attempts to connect to the first Terracotta Server Array server specified in the **URLs** list. The length of the pause is determined by the `watt.server.cachemanager.connectTimeout` parameter. If a connection is not established to any of the servers in the Terracotta Server Array within the time specified by this parameter, Integration Server logs an error and takes the action associated with the selected option for the **Action on Startup Error** parameter.

For more information about the **Action on Startup Error** parameter, see "[About Action on Startup Error Options](#)" on page 25.

For information about setting the `watt.server.cachemanager.connectTimeout` parameter, see *webMethods Integration Server Administrator's Guide*.

Solution

When Integration Server cannot connect to the Terracotta Server Array, do the following:

- Start the Terracotta Server Array if it is not already started.
- Make sure the machine on which Integration Server is running can reach the Terracotta Server Array host specified in the `tc-config.xml` file.
- Test the connection by pinging the servers in the Terracotta Server Array. The servers are listed in the **Terracotta Server Array URLs** list on the **Settings > Caching > Add Cache Manager** screen.
- Make sure the same version of Terracotta software is running on the client and the Terracotta Server Array. Check the `ehcache.log` file located in the

Integration Server_directory/instances/instance_name/logs directory. If the Terracotta software version numbers are not the same, a version mismatch error will be logged and Integration Server shuts down.

- Make sure the Integration Server has a valid Terracotta license key. If the Integration Server does not have a valid license key, you can enable clustering on the Integration Server; however, it will not be able to connect to the Terracotta Server Array. When the Integration Server starts up, it will write an error to the server log and then take the action associated with the selected option for the **Action on Startup Error** parameter. You can add a Terracotta license key by placing the Terracotta license file into the *SoftwareAG_directory \common\conf* directory of the Integration Server. You must restart the Integration Server after adding a Terracotta license. Alternatively, you can use the **Settings > License > License Details > Edit** page in Integration Server Administrator to point to a Terracotta license file at a different location. For more information about adding the Terracotta license file, see *webMethods Integration Server Administrator's Guide*.

Integration Server Cannot Connect to the Servers in the tc-config.xml File

If Integration Server has made the initial connection to the Terracotta Server Array and has obtained the tc-config.xml file, but the server cannot connect to any other servers in the Terracotta Server Array, Integration Server will wait indefinitely. Integration Server writes entries to the tc-client-logs and ehcache log files.

Solution

You can configure Integration Server to wait for a specified amount of time to connect to the Terracotta Server Array. After the wait time elapses, Integration Server takes the action associated with the selected option for the **Action on Startup Error** parameter. You configure the wait time for Integration Server by adding the following Java system properties to the custom_wrapper.conf file, which is located the *Software AG_directory/profiles/IS_instance_name/configuration* directory:

1. In a text editor, open custom_wrapper.conf.
2. Add two wrapper.java.additional.*n* parameters as follows:

```
wrapper.java.additional.n--Dcom.tc.ll.max.connect.retries=retryAttempts
wrapper.java.additional.n--Dcom.tc.ll.socket.reconnect.waitInterval= waitInterval_ms
```

Where:

- *n* is the next unused sequential number.
 - *retryAttempts* is the number of attempts Integration Server is to make to connect to the Terracotta Server Array.
 - *waitInterval_ms* is the number of milliseconds Integration Server is to wait before each attempt to retry the connection.
3. Save and close custom_wrapper.conf.

- Restart Integration Server for the changes to take effect. For more information about passing Java system properties to Integration Server, see *webMethods Integration Server Administrator's Guide*.

What Happens when Integration Server Is Disconnected from the Terracotta Server Array?

When an Integration Server in a cluster is disconnected from the Terracotta Server Array, the following occurs:

- Integration Server will accept new connections.
- Stateless sessions, new and existing, will continue to be fully functional. Note that stateless session information is not written to Terracotta Server Array.
- Stateful sessions, new and existing, will execute but will have limited functionality because stateless session information will not be replicated to other Integration Servers in the cluster as long as the Integration Server remains disconnected from the Terracotta Server Array. The Integration Server maintains its own session state information and does not have access to state information from other cluster members while disconnected from the Terracotta Server Array.

When executing a top-level stateful service, Integration Server attempts to write to the Terracotta Server Array after service invocation completes. If the **Timeout Behavior** is set to `localReads` or `exception`, the attempt to write to the Terracotta Server Array returns an exception. Integration Server writes this exception, often a `NonStopCacheException`, to the server log as part of the following error message:
[ISS.0033.0154E] Could not save session *sessionId* to the session cache.*exceptionMessage*

At the time Integration Server reconnects to the Terracotta Server Array, session state information becomes available across the cluster. However, all of the Integration Servers in the cluster might have different session state information. The first Integration Server to process a request updates the Terracotta Server Array with the new session state information. The updated session information becomes the session state for all of the Integration Servers in the cluster going forward. When another Integration Server begins to process a request for the session, that Integration Server retrieves the session state from the Terracotta Server Array. Consequently, session information on other Integration Servers in the cluster might be lost.

- How distributed caches in Integration Server behave during a loss of connection is controlled by the **Timeout**, **Immediate Timeout When Disconnected**, and **Timeout Behavior** settings for each distributed cache. For information about these settings, see *webMethods Integration Server Administrator's Guide*

Clients for a distributed cache continue to process requests, if all of the following are true:

- The **Timeout Behavior** parameter for each distributed cache is set to `localReads` or `noop`.

- The services do not attempt to write any data to a distributed cache for which the **Timeout Behavior** is localReads or exception.
- The list of cluster hosts on the **Settings > Cluster** page will not contain any cluster members.
- The Integration Server writes entries to the tc-client-logs and ehcache log files when it disconnects and reconnects to the Terracotta Server Array.

Solution

The way in which Integration Server rejoins the Terracotta Server Array depends on how long the connection is lost.

If...	Then...
The connection is momentarily interrupted	<p>If the Terracotta Server Array is configured to automatically reconnect, the Integration Server identity and state are retained on the Terracotta Server Array so that Integration Server can readily reconnect when the connection is restored. To learn more about the reconnect behavior of a Terracotta Server Array, see the information about automatic client reconnect in the <i>Configuring Terracotta Clusters For High Availability</i> documentation on the Terracotta website.</p> <p>If the Terracotta Server Array is not configured to automatically reconnect, the Integration Server will attempt to rejoin the Terracotta Server Array.</p>
The connection is lost for a long period of time	Integration Server attempts to rejoin the Terracotta Server Array. To prevent inconsistencies in your distributed cache, Terracotta automatically acquires a new cache state from the Terracotta server after rejoining the Terracotta Server Array. For more information about the rejoin and nonstop behavior of Integration Server, see <i>webMethods Integration Server Administrator's Guide</i> .

What Happens When Terracotta Servers in the Terracotta Server Array Are Disconnected?

If the network connection is interrupted between the active server and one or more of the standby servers in the Terracotta Server Array, a situation might arise where one of the standby servers becomes an active server. This situation could result in a split brain

scenario. In a split brain scenario, the connection between the servers is restored but two or more Terracotta servers assume the role of the active server. When this happens, the data shared between the servers can become lost or corrupted.

To prevent a split brain scenario, Terracotta determines which server should become the new active server based on how many clients were connected to each server.

How Do Integration Servers Respond when Terracotta Servers in the Terracotta Server Array Are Disconnected?

If the network connection is interrupted between the active server and one or more of the standby servers in the Terracotta Server Array, the Integration Servers that were connected to the active server will continue to use the same active server. However, if a new Integration Server joins the cluster, it will attempt to join the first Terracotta Server defined in the **Terracotta Server ArrayURLs** list. With this behavior, it is possible that new Integration Servers joining the cluster will connect to a new active server. This situation can lead to data loss when the servers reconnect. Additionally, while both servers are split, the Integration Server cluster might behave unpredictably because the Integration Servers in the cluster would have access to different cached data.

To avoid this situation, Software AG recommends that you shut down the new active server until the connection between the Terracotta servers is restored.

Solution

To prevent these issues, make sure the network connection between active server and standby servers in the Terracotta Server Array is reliable. For more information about preventing a split brain scenario, see "Split Brain Scenario" in the Terracotta Server Arrays Architecture documentation on the Terracotta website.

Index

A

- Access Control Lists (ACLs) for clustered servers 19
- adding
 - server back into cluster 33
 - servers to cluster 22
- audit database
 - special considerations for use in a cluster 19

B

- benefits
 - of clustering 8
 - of failover support 9
- Broker
 - special considerations for use in a cluster 19

C

- client applications
 - how client code redirects 15
- cluster database
 - session objects, description 13
- clustering
 - benefits 8
 - benefits of using with guaranteed delivery 28
 - description 8
 - requirements before setting up 18
 - setting up server environments 18
 - software requirements 18

D

- database, cluster
 - session objects, description 13
- deleting
 - Integration Server from cluster 32
- disabling server clustering 32
- documentation
 - using effectively 5

E

- environments, server
 - setting up for clustered servers 18
- event subscriptions for clustered servers 19
- eventcfg.bin file 19

F

- failover support
 - benefits 9
 - description 9
- firewall
 - specifying IP domain name and address for use outside of 23

G

- guaranteed delivery
 - benefits of using with clustering 28
 - description 10
 - overview 10

I

- installing
 - Integration Server 18
- Integration Servers
 - Access Control Lists (ACLs) for clustered servers 19
 - adding server back to cluster 33
 - event subscriptions for clustered servers 19
 - packages on clustered servers 18
 - removing server from your cluster 32
 - software requirement for server clustering 18
 - taking offline 33
 - user accounts on clustered servers 19
- IP domain name and address
 - specifying for external use 23
- IS Servers
 - setting up server clustering 18
 - setting up server environments for clustering 18

J

- job store 28
- job store locking 28

L

- load balancing
 - overview 9
- locking
 - job store 28

M

- messaging provider, considerations for use in a cluster 19

N

Node Identity configuration setting
specifying 23

O

overview
failover support 9
guaranteed delivery 10

P

packages on clustered servers 18

R

redirecting requests
how clients redirect 15
reinstating server into cluster 33
removing server from cluster 32
requirements
software requirements for server clustering 18

S

server clustering
Access Control Lists (ACLs) for clustered servers
19
adding new Integration Servers 22
adding server back to cluster 33
before you set up 18
disabling 32
event subscriptions for clustered servers 19
packages on clustered servers 18
removing Integration Server from cluster 32
session objects description 13
setting up on server 18
software requirements 18
specifying nodes for external use 23
taking a server offline 33
user accounts on clustered servers 19
viewing
servers in cluster 32
session objects
description 13
software requirements for server clustering 18

T

Terracotta
enabling clustering with 22
setting up the Terracotta Server Array 21
Terracotta Server Array URLs 20
using as the caching software 21

U

Universal Messaging
special considerations for use in a cluster 19
user accounts on clustered servers 19

V

viewing
servers in cluster 32

W

webMethods guaranteed delivery
description 10