

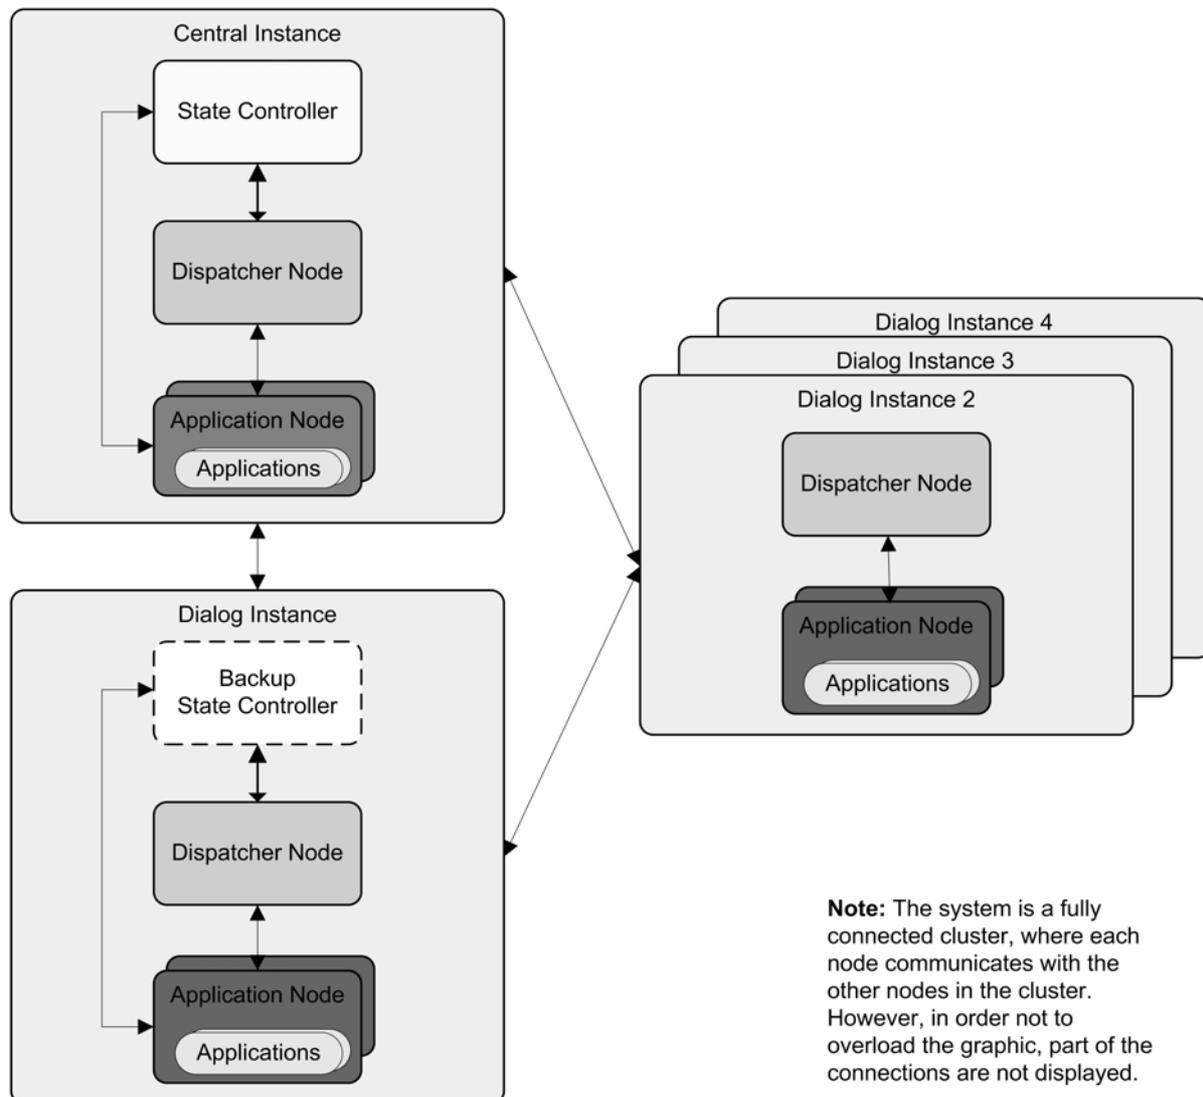
The Revised SAP J2EE Engine 6.20 Cluster Architecture

Overview

Since PL21 the SAP J2EE Engine 6.20 has introduced a cluster architecture that considerably improves overall stability and reliability of the cluster. This document explains the revised cluster architecture and describes the configuration procedures for installing a new SAP J2EE Engine 6.20 and for upgrading a previous version to PL21 or higher.

The revised SAP J2EE Engine 6.20 cluster architecture consists of:

- One state controller that manages the state of the whole cluster – keeping locks, configurations, and so on. Optionally, to avoid single point of failure, there can also be another, backup state controller.
- One or more dispatcher nodes to dispatch requests to the application nodes
- One or more application nodes that handle client requests



Note: The system is a fully connected cluster, where each node communicates with the other nodes in the cluster. However, in order not to overload the graphic, part of the connections are not displayed.

State Controller and Backup State Controller

The new feature here is the state controller and the centralized management of the cluster. There must be exactly one state controller in the cluster. Optionally, there can be one backup state controller that takes on responsibility for the cluster if the original state controller fails. If the cluster is configured to have more than one state controller, the Cluster Manager of the state controller will not start the additional state controllers.

The state controller serves as a single, central, persistent data storage area in the cluster and does not handle client requests. It also holds global information about the cluster and the participating nodes – for example, cluster-wide locks, and so on.

If the state controller fails, the backup state controller takes on the role of the state controller. When the original state controller has been restarted, it becomes the new backup state controller.

If the backup state controller fails while the state controller is still running, the cluster will still function correctly.

To prevent cluster failures in the case of hardware problems, we recommend that the two state controllers reside on different physical machines.

Mechanisms for Automatic Restarts in Case of Cluster Node Failures

The 6.20 cluster architecture is designed so as if a cluster node fails, it automatically restarts. This guarantees the stability of the whole system and ensures that the deployed applications will continue to function correctly regardless of cluster node failures. The behavior of the cluster in case of failures is as follows:

- If an application node fails, it automatically restarts
- If the state controller fails and there is no backup state controller, all nodes in the cluster automatically restart.
- If the state controller fails and there is a backup state controller installed, the cluster continues to function correctly. The backup state controller becomes the state controller. The original state controller automatically restarts and joins the cluster as a backup state controller.
- If the backup state controller fails, it automatically restarts. As the state controller is still running, the cluster will still function correctly.
- If the dispatcher node fails, it automatically restarts.

Note: These mechanisms work when you use the cyclic scripts to start the cluster nodes, when you use the startup framework, and when you are in an ABAP environment. If you use the standard *go* scripts to start the nodes, these mechanisms for automatic restart of the cluster nodes will not be available.

Cluster Startup

When starting the cluster nodes, follow the steps below:

1. Start the state controller first.
2. Start the remaining dispatcher and application nodes in the cluster.
These nodes will connect to the state controller.
3. Start the backup state controller.
It will connect to the state controller. The existence or successful start of the backup state controller is not critical for successful start of the whole cluster.

Note: To initially start the cluster, you always have to start the state controller first. You cannot start the cluster by using the backup state controller as a state controller. The backup state controller can only take the responsibilities of the state controller in an already started and running cluster when the state controller has failed, but the backup state controller cannot be used to start a new cluster.

For Enterprise Portal customers we recommend that after you install the SAP J2EE Engine, you install the startup framework and use it to start the SAP J2EE Engine. The startup framework provides centralized management of the cluster nodes and monitors their life cycle. In case of cluster node failure, the framework automatically restarts the corresponding node. For more information about the startup framework, see the *Installing and Configuring the 6.40 Startup Framework to Use with SAP J2EE Engine 6.20* guide attached to SAP Note 748713.

Note: For more information about the different startup options, see the *Installation and Upgrade Manual*.

Configuration Settings for the Cluster Nodes

General Settings

To ensure correct communication and operation within the cluster, the following properties must be configured in the Cluster Manager of each cluster node:

- *ClusterHosts* – specifies the set of hosts and ports to which the corresponding node tries to connect to join the cluster. See the *Installing a New Cluster* section below to determine when the system sets the value of this property automatically, and when you have to configure it manually.
- *ClusterElementType* – determines the type of cluster node; possible values: *StateController*, *BackupStateController*, *Dispatcher*, *ApplicationNode*. This property is automatically set by the system, when you create the corresponding node.

If the settings for the *ClusterHosts* and *ClusterElementType* properties of a cluster node are not correct, this node will not be able to start and join the cluster. Thus, if the cluster is configured to have more than one state controller, only one of these nodes will be started. If two backup state controllers are configured, but no state controller, the cluster will not be started.

The *SuspendedElement* property in the Service Manager must be *true* for the state controller and the backup state controller, and *false* for the application nodes. This property is also configured by default when you install the cluster nodes.

Changed Behavior and Configuration Tuning

If you upgrade an existing cluster configuration, you have to check the following configurations:

- The *-server* option must not exist in the Java VM type settings for SAP J2EE Engine 6.20 PL21 or higher. Therefore, if during the upgrade procedure you have received a warning, make sure that you do not have a *-server* option in your Java VM settings. Otherwise, you will not be able to run the SAP J2EE Engine.
- Check the value of the *AdditionalLoadTimeout* property in the Service Manager. In PL23 and higher this property is used as a watchdog and if its value is small, the nodes might not be able to start and connect to the cluster.
- Check the values of the *PingInterval* and the *PingReplyTimeout* properties in the Cluster Manager. These properties are introduced in PL22 and are used to detect network disconnection in the cluster. As they are enabled by default, make sure you configure them according to your cluster configuration.

Additional Cluster Configuration Settings

You can also configure the following options:

- Configure a dispatcher node to distribute requests only to the local application nodes (use the *LocalLoadBalancing* property of the Service Manager on the dispatcher node)
- Tune the cluster nodes to detect network disconnection (use the *PingInterval*, *PingReplyTimeout*, and *ClusterElementIPAddress* properties of the Cluster Manager)
- Detect inactivity or failures of cluster nodes – this includes configuring the internal watchdogs that cope with occasional hang-ups and prevent the nodes from being stuck
- Configure the cluster to restart the cluster nodes after they have processed a certain number of requests and notify the system administrator by e-mail (use the Ageing Service)

Note: For more information about configuring a SAP J2EE Engine 6.20 cluster, see the *Administration Manual* → *Administration of SAP J2EE Engine 6.20 Cluster*.

Installing a New Cluster

This section describes the procedure for installing a new SAP J2EE Engine 6.20 cluster – that is, when there is no previous version installed.

Default Installation

The default installation procedure installs a cluster consisting of one state controller, one dispatcher node, and one application node. Note that the default installation procedure for earlier versions of SAP J2EE Engine 6.20 – that is, prior to PL21 – installed a cluster consisting of one dispatcher and one server node.

The installation procedure configures the state controller by setting the *ClusterHosts* property in the Cluster Manager to empty. It configures the dispatcher and application nodes by setting their *ClusterHosts* properties to contain the host and port for the state controller. Thus, the dispatcher and the application nodes are configured to connect to the state controller. The procedure also configures the *ClusterElementType* property for all cluster nodes.

Adding a Dialog Instance to the Cluster

A dialog instance is a box in the cluster (a physical machine with all dispatcher and application nodes running on it) that does not contain a state controller and connects to a central instance (a box that contains a state controller).

To add to the cluster a dialog instance without a backup state controller, use the standard installation procedure but do not install a state controller or a backup state controller. During the installation, specify the hosts and the ports of the state controller and the backup state controller in the cluster. The installation procedure will then automatically configure the correct *ClusterHosts* property values for the cluster nodes you are installing.

Installing a Backup State Controller

During a dialog instance installation, select the *Create Backup Controller* option and specify the host and port for the state controller. The installation procedure will then automatically configure the correct *ClusterHosts* property values for the cluster nodes you are installing.

Note: Use the Config Tool to manually specify the host and port of the backup state controller in the *ClusterHosts* property of all cluster nodes already installed since this will not be configured automatically.

Adding Dispatcher or Application Nodes to an Existing Cluster Box

To add additional nodes to the cluster, use the Config Tool. It automatically configures the values of all necessary properties (*ClusterHosts* and *ClusterElementType*).

Maximum Number of Cluster Nodes

For technical reasons, the maximum number of cluster nodes in a cluster can be 32. This includes state controller, backup state controller, dispatcher nodes, and application nodes.

Note: For more information about the system requirements and the installation procedures on different operation systems, refer to the *Installation and Upgrade Manual*.

Note: When you add nodes to an existing cluster, which is already configured to work with the startup framework, make sure you also configure the new node to use the startup framework, as you have already configured the other nodes in the cluster. Otherwise, the life cycle of the new node will not be managed by the startup framework.

Upgrading an Existing Cluster

This section describes what the installation procedure does when you upgrade a previous version of SAP J2EE Engine 6.20 to PL24.

The upgrade procedure does not differ from the installation in terms of user input. The differences refer to the background activities, which the upgrade procedure performs automatically.

The installation procedure:

- Configures all server nodes as application nodes.
- Configures the dispatcher nodes.

- Creates and configures a state controller or a backup state controller if you have selected to install these. If you upgrade from SAP J2EE Engine PL21 or PL22, the procedure configures the existing state controller and backup state controller.

Install the state controller on a host where you have a primary server.

Make sure you have installed only one state controller in the whole cluster. If you have installed more than one state controller, use the Config Tool to delete the unnecessary state controllers.

To add additional nodes to the cluster after you have upgraded your current installation, follow the appropriate steps described in the *Installing a New Cluster* section above.

Upgrading a Cluster Which Uses the Startup Framework

If you upgrade SAP J2EE Engine PL 19 or 20, the upgrade procedure adds a new node (the state controller or the backup state controller) to the existing cluster configuration. Therefore, you have to manually configure that node to use the startup framework. Otherwise, the life cycle of the new node will not be managed by the startup framework.

Note: For more information about the upgrade procedure, refer to the *Installation Manual* → *Upgrading SAP J2EE Engine 6.20*.

Note: For the necessary steps, when adding new cluster nodes to an already upgraded installation of the SAP J2EE Engine 6.20 cluster with an Enterprise Portal 6.0 running on it, refer to the *Installing an Additional SAP J2EE Engine for a Portal Cluster* guide included in the documentation set *EP 6.0 Installation Guide* (zip file), available on SAP Service Marketplace at: service.sap.com/ep60 → *Documentation & More* → *EP 6.0 (SP2 and before)* → *Installation*.