

Clustering and farming

These topics are about how to set up a cluster and perform server farming.

A *cluster* or *server farm* is a group of servers that transparently provide enterprise services, such as Servlets and JavaServer Pages (JSP) support, as if it was a single server. The servers, typically running on separate systems, exchange messages to synchronize data, allowing any individual node to process requests for a distributed application and to take over a user's session when its node fails. Configuring multiple servers into a cluster is commonly called *clustering*.

Consider using a cluster of servers when you need to improve the scalability or the availability of your Web applications.

- Clustering improves scalability by allowing you to increase the system capacity by adding extra servers. Load balancing is achieved by adding a component, like the remote Apache HTTP Server with the Apache Tomcat JK module that sprays requests to each node in the cluster.
 - Running multiple nodes on the same host, called *vertical clustering*, is used when limitations in a single node prevent it from fully utilizing the capacity of the host.
 - Running multiple nodes on different hosts, called *horizontal clustering*, is used to add capacity by adding an additional host. Horizontal clustering is more common than vertical clustering because it also improves availability.
- Clustering improves availability by automatically synchronizing state data among all the nodes in the cluster. If any node in the cluster fails, subsequent requests can be processed by any other node in the cluster. This process, typically called *failover*, occurs on subsequent requests. Any requests being processed by a node when it fails will still be lost.

This section is organized as follows:

- [EJB failover](#) — In Geronimo, you can configure OpenEJB for failover scenarios <http://openejb.apache.org/3.0/failover.html>, such as multipoint and multicast. Further more, Geronimo uses WADI to support failover for Stateful Session Bean (SFSB) in a unicast way.
- [Farming using Deployment](#)
- [JMS clustering in Geronimo](#) — JMS clustering in Geronimo is handled by ActiveMQ component directly.
- [Plugin based Farming](#)
- [Tomcat Native Clustering](#) — The Tomcat Web container provides a native clustering solution that can be configured through Geronimo using gbean definitions within config.xml or using server.xml starting from v2.2 just like you did for a standalone Tomcat server.
- [WADI Clustering](#)