Sun Cluster 2.2 7/00 Data Services Update: Apache Web Server

Please
Recycle

Adobe PostScript

Sun Cluster 2.2 7/00 Data Services Update: Apache Web Server

Sun™ Cluster High Availability (HA) for Apache is an Apache Web Server that runs under the control of Sun Cluster. This document describes how to install the Apache Web Server so that it runs as the Sun Cluster HA for Apache Web Server data service.

Preparing for Installation

The Apache Web Server requires some variations from the default installation parameters. In particular:

- You must specify the logical host name, *not* the physical host name.
- Your binary files can reside on the local host, in which case you may not have to take down the logical host for upgrades; however, the `docs` directory must reside on the multihost disk. Alternatively, you can place all the files, including the binaries, on the multihost disk. In the latter case, when prompted, do *not* use the default server root disk.

The procedures described in this document assume that you are installing the binaries on the multihost disk.

You must follow certain conventions when you configure URL mappings for the Web server. For example, to preserve availability while setting the CGI directory, you must locate the mapped directories on the multihost disks associated with the logical host that serves Apache Web Server requests for this mapping. In this case, you map your CGI directory to
`/`*logicalhost*`/commerce/apache-home/cgi-bin`.

In situations where the CGI programs access back-end data, verify that the data are also located on the multihost disks associated with the logical host that serves the Apache Web Server requests.

In situations where the CGI programs access back-end servers, such as an RDBMS, ensure that the back-end server is also controlled by Sun Cluster. If the server is an RDBMS supported by Sun Cluster, use one of the highly available RDBMS packages, such as Sun Cluster HA for Oracle, Sun Cluster HA for Sybase, or Sun Cluster HA for Informix. Otherwise, you can use the APIs that are documented in the *Sun Cluster 2.2 API Developer's Guide* to put the server under Sun Cluster control.

Installing the Apache Web Server Software

For standard Apache installation instructions and details on the Apache Web Server software, see the Apache Web site at:

```
http://www.apache.org
```

▼ How to Download, Uncompress, and Untar the Binaries

Sun Cluster 2.2 supports Apache Web Server version 1.3.9 for the Solaris™ 8 operating environment. You can download this version from the Apache Web site at `http://www.apache.org` either as a binary distribution or as source files. If you choose the source files, you can compile them on the local machine. We recommend the binary distribution. Do the following:

1. **Make a directory for the Apache installation script:**

   ```
   % mkdir script-dir
   ```

2. **Go to the Apache Web site:** `http://www.apache.org` **and download the binary distribution. Save the file in** *script-dir***.**

   Typically, the downloaded file is in compressed `gzip` and tarred format named *filename*`.gz.tar`.

3. **Uncompress the binaries by typing:**

   ```
   % gzip -d filename.gz.tar
   ```

4. **Untar the resulting file by typing:**

   ```
   % tar -xvf filename.tar
   ```

   The directory then contains the installation scripts.

▼ How to Install the Apache Web Server

1. **Run the** `install-bindist.sh` **script and supply the name of the directory on the multihost disk on which the server is to be installed, as follows:**

   ```
   phys-hahost1# ./install-bindist.sh install-dir
   ```

   The *install-dir* directory is the multihost disk on which the binary distribution is to be installed. If you do not specify *install-dir*, Apache uses `/usr/local/apache` as the destination directory, which resides on the local disk only.

   The following is then displayed:

   ```
   install-bindist.sh install-dir
   Installing binary distribution for platform sun4u-sun-solaris2.80
   into directory install-dir …
   Ready.
   You now have successfully installed the Apache 1.3.9 HTTP server.
   To verify that Apache actually works correctly you should first
   check the (initially created or preserved) configuration files:
   install-dir/conf/httpd.conf
   You should then be able to immediately fire up Apache the first
   time by running:
   install-dir/bin/apachectl start
   Thanks for using Apache.
   The Apache Group
   http://www.apache.org/
   ```

2. **Edit the** *install-dir*/`conf/httpd.conf` **file, as follows:**

   a. **Redefine the property of** `ServerName` **to be the logical host name (rather than the physical host name).**

      This step is necessary for Sun Cluster HA for Apache to fail over correctly.

   b. **Set the property port to the port on which the Apache Web Server listens.**

      The default and most commonly used port number for Apache is 80.

3. **(Optional) Further customize** *install-dir*/`conf/httpd.conf` **as required.**

Installing and Configuring Sun Cluster HA for Apache

Prior to configuring Sun Cluster HA for Apache, you must install and set up the Sun Cluster and Apache products.

You configure Sun Cluster HA for Apache by using the `hadsconfig`(1M) command. See the man page for details.

▼ How to Install and Configure Sun Cluster HA for Apache

1. **On all nodes, install the Sun Cluster patch that enhances the** `scinstall`**(1M) command to recognize Sun Cluster HA for Apache. The patch is available from your service provider or from the Sun patch Web site** `http://sunsolve.sun.com`**.**

2. **On each Sun Cluster node that is a potential master of the logical host on which Apache is installed, run the** `scinstall`**(1M) command to install the Sun Cluster HA for Apache package,** `SUNWscapc`**, and supporting packages in the default location.**

3. **On all nodes, install the Sun Cluster patch that enhances the** `hadsconfig`**(1M) command to recognize Sun Cluster HA for Apache.**

4. **On one node only, run the** `hadsconfig`**(1M) command to create, edit, or delete instances of the Sun Cluster HA for Apache, as follows:**

   ```
   phys-hahost1# hadsconfig
   ```

   For information on input to supply to `hadsconfig`(1M), see "Understanding the Configuration Parameters" on page 5.

5. **If the service is not yet registered, run the** `hareg`**(1M) command on one node only to register the Sun Cluster HA for Apache data service.**

   To register the service on the logical host only, include the -h option and logical host name, as follows:

   ```
   phys-hahost1# hareg -s -r apache -h logicalhost
   ```

6. **On one node only, run the** `hareg`**(1M) command to enable Sun Cluster HA for Apache and perform a cluster reconfiguration, as follows:**

```
phys-hahost1# hareg -y apache
```

7. **Confirm your Apache Web Server operation by establishing a** `telnet` **connection to your logical host, where the service is listening. For example:**

```
# telnet logicalhost port-number
```

Understanding the Configuration Parameters

This section describes the information you supply to the hadsconfig(1M) command to create configuration files for Sun Cluster HA for the Apache data service. The hadsconfig command uses templates to create these configuration files. The templates contain defaults; some are hard coded and others are unspecified parameters. You must provide values for the unspecified parameters.

The fault probe parameters, in particular, can affect the performance of Sun Cluster HA for Apache data services. Tuning the probe interval value too low (that is, increasing the frequency of fault probes) might encumber system performance and result in unnecessary takeovers or attempted restarts when the system is simply slow.

Fault monitor parameters are configurable for Sun Cluster HA for the Apache Web Server. One of the parameters is the takeover flag, which specifies how Sun Cluster handles partial failover. You have two options:

- **y (yes)** — Sun Cluster attempts to switch over the logical host to another master. However, if the attempt fails, the logical host remains on the original master. This is the default setting.
- **n (no)** — Sun Cluster does not move the logical host to another master even if it detects problems with the data server, nor does it take any action against the faulty data server or database on the logical host.

Configure the Sun Cluster HA for Apache Web Server parameters listed in the hadsconfig(1M) input form by supplying the options described in TABLE 1-1. The parameters are listed in the order they appear in the form.

TABLE 1-1   Configuration Parameters for Sun Cluster HA for Apache

Parameter	Description
NAME (name of the instance)	The name tag that serves as an identifier for the instance. The log messages generated by Sun Cluster refer to this name tag. The hadsconfig(1M) command prefixes the package name to the value you supply here. For example, if you specify apache_80, hadsconfig produces SUNWscapc_apache_80.
LOGICAL_HOST (logical host)	The name of the logical host that provides service for this instance of Sun Cluster HA for Apache Web Server.
PORT (server port number)	The unique port for this instance of Sun Cluster HA for Apache Web Server. This is the server port value you supplied to the httpd.conf file while configuring Apache Web Server.

TABLE 1-1   Configuration Parameters for Sun Cluster HA for Apache *(Continued)*

Parameter	Description
INTERVAL (probe interval)	The time, in seconds, between fault probes. The default interval is 60 seconds.
TIMEOUT (probe timeout)	The time, in seconds, after which a fault probe times out. The default timeout value is 30 seconds.
TAKEOVER (takeover flag)	The flag that specifies whether a failure of this instance causes a takeover or failover of the logical host associated with the data service instance. Possible values are y (yes) and n (no).
BIN_DIR (binary directory)	The directory of the product installation that contains the httpd binary, for example, /hahost1/apache-home/bin.
CONF_DIR (configuration directory)	The directory of the product installation that contains the configuration file httpd.conf, for example, /hahost1/apache-home/conf.