Amigopod

High Availability Deployment Guide

ARUBA networks

Technical Note

Copyright

Open Source Code

Certain Aruba products include Open Source software code developed by third parties, including software code subject to the GNU General Public License (GPL), GNU Lesser General Public License (LGPL), or other Open Source Licenses. The Open Source code used can be found at this site:

http://www.arubanetworks.com/open_source

Legal Notice

The use of Aruba Networks, Inc. switching platforms and software, by all individuals or corporations, to terminate other vendors' VPN client devices constitutes complete acceptance of liability by that individual or corporation for this action and indemnifies, in full, Aruba Networks, Inc. from any and all legal actions that might be taken against it with respect to infringement of copyright on behalf of those vendors.

Warranty

This hardware product is protected by the standard Aruba warranty of one year parts/labor. For more information, refer to the ARUBACARE SERVICE AND SUPPORT TERMS AND CONDITIONS.

Altering this device (such as painting it) voids the warranty.

ARUBA®
networks

www.arubanetworks.com
1344 Crossman Avenue
Sunnyvale, California 94089
Phone: 408.227.4500
Fax 408.227.4550

Table of Contents

1 Introduction

This technical note describes the process of setting up a high-availability cluster of Amigopod Visitor Management Appliances.

Audience

This document is intended for network administrators and system integrators deploying an Amigopod-based visitor management solution.

Basic familiarity with the Amigopod Visitor Management Appliance is assumed. For in-depth information about the features and functions of the Amigopod appliance, refer to the Amigopod Deployment Guide.

Document Overview

The first section of the document explains the architecture of an Amigopod high availability system.

The next section contains a walkthrough and configuration guide for deploying a cluster of Amigopod Visitor Management Appliances.

2 About High Availability

Terminology & concepts

A **cluster** consists of a primary node and a secondary node, configured so that a failure of either node will not prevent the cluster as a whole from performing its normal functions.

The **primary node** is the active server in a cluster. The cluster's network services are always delivered by the primary node.

The **secondary node** is the backup server in a cluster. If the primary node fails, the secondary automatically takes over and continues delivering network service.

Fault tolerance is the ability of a server cluster to continue operating if either the primary or secondary node experiences a hardware failure.

Fail-over is the process by which the secondary node assumes control of the cluster once the primary node has failed.

A cluster's **virtual IP address** is a unique IP address that will always be assigned to the primary node of the cluster. In order to take advantage of the cluster's fault tolerance, all clients that use the cluster must use the cluster's virtual IP address, rather than each node's IP address.

Replication is the process of ensuring that the secondary node maintains an exact copy of the primary node's database contents and configuration. Replication is used to ensure that if a fail-over is required, the secondary node can continue to deliver an uninterrupted service to clients of the cluster.

Keep-alive is the process by which cluster failures are detected. The primary and secondary nodes verify that each is able to communicate with the other node by sending network requests and answering with a response.

Database replication is the process of ensuring that all changes to the database, including new guest accounts, changes to existing guest accounts, RADIUS roles, NAS servers, and RADIUS accounting information, are replicated from the primary node to the secondary node. This replication process occurs continuously in a normally operating cluster. Replication is required so that in the event of a primary node failure, the secondary node is up to date and can continue to deliver the same network services to clients.

Configuration replication is similar to database replication, but occurs at a slower rate due to the reduced frequency of configuration updates.

The **downtime threshold** is the time for which the primary node of the cluster must remain offline before an automatic fail-over will be initiated. This is 30 seconds by default.

Network architecture

Diagram 1 shows the network architecture for a high availability cluster.

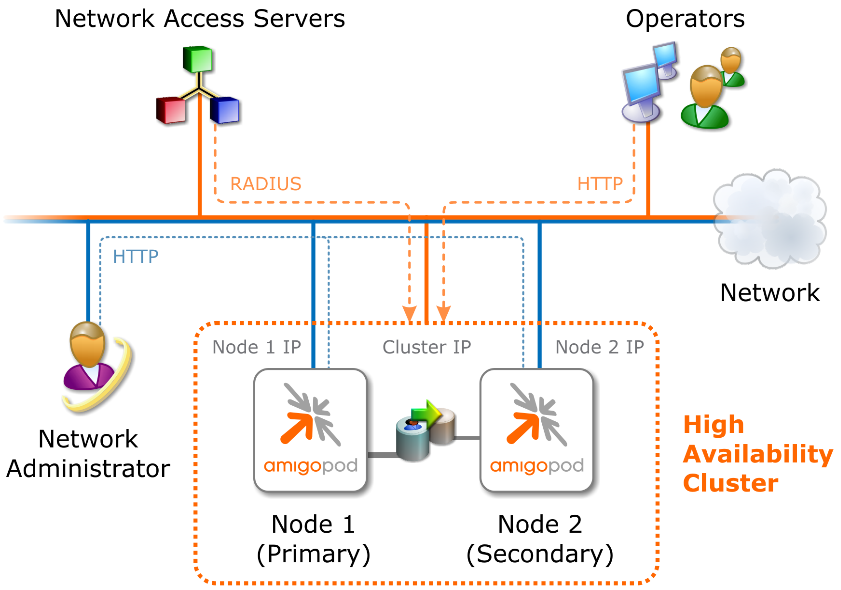

Diagram 1: Network architecture of high availability cluster

The key points to note about this architecture are:

- The RADIUS and web server protocols (HTTP and HTTPS) are supported by the cluster.
- The cluster has three IP addresses: each node has its own IP address, and there is a virtual IP address for the cluster which will always be assigned to the primary node in the cluster.
- For the cluster to provide fail-over redundancy, all network access servers and operators must use the cluster's IP address.
- The network administrator should use the node IP addresses to perform system administration tasks on each node, including managing the cluster itself.

The nodes in the cluster must be connected to the same local network. Use high quality network cables and reliable switching equipment to ensure the nodes have an uninterrupted network connection.

NOTE There should be no routers, gateways, firewalls, or network address translation (NAT) between the two nodes.

NOTE Having nodes in different physical locations is not recommended and is not a supported configuration for the cluster.

Operating a High Availability cluster

Normal cluster operation

When the cluster is operating normally, the cluster status will be:

✓ The cluster is running normally.

In this state, the primary node is assigned the cluster IP address and is responsible for delivering network services to clients. Each node is also continuously performing failure detection, database replication and configuration replication, as explained below.

Failure detection

Failure detection is accomplished using a **keep-alive test**. The primary and secondary nodes verify that each is able to communicate with the other node by sending network requests and answering with a response. This takes place at the Keep Alive Rate specified in the cluster configuration, which by default is once every 2 seconds.

If several consecutive keep-alive tests have failed, the cluster determines that a failure has occurred. A cluster fail-over may then take place, depending on which node has failed. Refer to section 0 for information about a primary node failure, or section 0 for information about a secondary node failure.

To avoid any network service interruptions, it is important that the nodes maintain an uninterrupted network connection.

Database replication

Database replication occurs continuously in a normally operating cluster. All database modifications, including new guest accounts, changes to existing guest accounts, RADIUS roles, NAS servers, and RADIUS accounting information, are replicated from the primary node to the secondary node. The replication delay will depend on the volume of database updates and system load but is generally only a few seconds.

Replicating the database contents ensures that in the event of a primary node failure, the secondary node is up to date and can continue to deliver the same network services to clients.

NOTE While the primary node is online, the secondary node's database can only be updated with replication changes from the primary node. No other database changes can take place on the secondary node. Because of this, any form that requires a database update will be disabled and shown as "Read Only Access" on the secondary node.

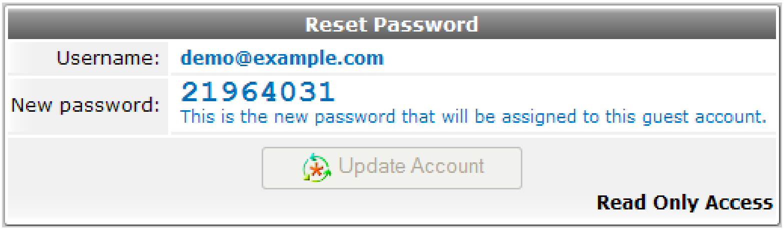

Reset Password

Username: demo@example.com

New password: 21964031
This is the new password that will be assigned to this guest account.

Read Only Access

Ensure that you always access the cluster using the virtual IP address when performing any database updates, such as creating new guest accounts or performing RADIUS authentication. This is required so that the changes will be performed on the primary node and then replicated to the secondary node.

Configuration replication

Configuration replication also occurs continuously within the cluster, but takes place at a slower rate due to the reduced frequency of configuration updates. This rate is the Config Sync rate specified in the cluster configuration, which by default is once every minute.

The configuration items that are replicated include:

- Configuration for installed plugins
- Fields defined in Guest Manager
- Forms and views defined in Guest Manager
- Guest self-registration pages
- Instances of reports that have previously been run
- LDAP authentication servers and translation rules
- Network login access configuration
- Operator login configuration
- Operator logins
- Operator profiles
- Print templates defined in Guest Manager
- Publicly-accessible web server items in Content Manager
- RADIUS server configuration
- Report definitions
- SMS service configuration
- SMTP server configuration
- SMTP settings for email receipts
- SNMP server settings
- The set of currently installed plugins
- Web Login pages

Certain configuration items are not replicated. These are:

- HTTP Proxy settings
- Network interface configuration
- RADIUS dictionary entries
- SSL certificate settings
- Subscription IDs in Plugin Manager
- System hostname

Primary node failure

If the cluster's primary node fails, the cluster status will be displayed on the secondary node as:

The secondary node is running, but the primary node is down or stopped.

While the primary node is down, the cluster is in a failed state and cannot deliver network services. If the primary node recovers within the downtime threshold, the cluster will automatically return to the normal state and network service will be restored.

An automatic fail-over will be initiated after the primary node has been offline for the **downtime threshold**, which is 30 seconds by default.

Once fail-over has occurred, the cluster status will be displayed on the secondary node as:

The secondary node has taken over the cluster services because the primary node is down.

In the fail-over state, the secondary node will assume control of the cluster and will take over the cluster's IP address. This will restore network service for clients of the cluster. Replication will stop as there is no longer a primary node.

While the primary node is offline, the cluster will no longer be fault-tolerant. A subsequent failure of the secondary node will leave the cluster inoperable.

Refer to section 0 for instructions on recovering a cluster in this state.

The secondary node has taken over the cluster services. The primary node is back online, but the cluster needs to be recovered.

In this state, the primary node was offline for a period of time greater than the downtime threshold, and then recovered. The cluster has failed over to the secondary node.

In this state, the cluster is not fault-tolerant. A subsequent failure of the secondary node will leave the cluster inoperable.

Recovering the cluster is required for replication to resume and return the cluster to a fault-tolerant state.

Refer to section 0 for instructions on recovering a cluster in this state.

Secondary node failure

If the cluster's secondary node fails, the cluster status will be displayed on the primary node as:

The primary node is running, but the secondary node is down or stopped.

The cluster will continue operating without service interruption. Network services will be unaffected as the cluster's virtual IP address is assigned to the primary node.

While the secondary node is offline, the cluster will no longer be fault-tolerant. A subsequent failure of the primary node will leave the cluster inoperable.

To recover the cluster, the secondary node must be brought back online. If the node has experienced only a temporary outage and has the same cluster configuration, the cluster will automatically repair itself. Replication will update the secondary node with any database or configuration changes that were made on the primary node while the secondary node was offline.

If the secondary node was replaced due to a hardware failure then the cluster must be destroyed and rebuilt. Refer to section 0 for instructions on recovering a cluster in this state.

Cluster status

The current status of the cluster is shown at the top of each page that is related to High Availability Services.

Refer to this table for an explanation of each possible status, and the recommended action to take, if any.

Status	Description
⊖	This system is not part of a high availability cluster. ▪ To create a new cluster and make this server the primary node, use the **Create New Cluster** command. ▪ To join a cluster and make this server the secondary node, use the **Join Cluster** command.
✓	The cluster is running normally. Click the ℹ View details link to show more information about the cluster. ▪ To perform a scheduled maintenance task, such as a reboot, on the primary node in the cluster, use the **Cluster Maintenance** command. ▪ Refer to section 0 for more information about normal cluster operations.
✗	The secondary node has taken over the cluster services because the primary node is down. ▪ A fail-over has occurred. The cluster must be recovered to resume fault-tolerant operation. ▪ Ensure the primary node is back online.
✗	The secondary node has taken over the cluster services. The primary node is back online, but the cluster needs to be recovered. ▪ A fail-over has occurred. The cluster must be recovered to resume fault-tolerant operation. ▪ Use the procedure described in section 0.
⊖	A failure has occurred. ▪ Check the detailed status information. ▪ If this message persists, you may need to rebuild the cluster. Refer to section 0.
⚠	The primary node is running, but the secondary node is down or stopped. ▪ The secondary is no longer available. Check the Remote Status on the primary node to determine the cause of the problem. ▪ To clear the error condition, bring the secondary node back online. The cluster will return to fault-tolerant mode automatically. ▪ If the secondary node needs to be replaced, the cluster must be rebuilt. Refer to section 0.
⚠	The secondary node is running, but the primary node is down or stopped. ▪ The primary is no longer available. Check the Remote Status on the secondary node to determine the cause of the problem. ▪ The cluster IP address is inaccessible and network services are unavailable. ▪ Automatic fail-over will take place after the downtime threshold has been exceeded.
⚠	The cluster services are starting. Check the detailed status information.

⚠	The primary node is running, but a problem has been detected. Check the detailed status information.
⚠	The primary node is running, but the secondary node is reporting a problem. Check the detailed status information.
ℹ	The cluster is recovering from a failure. Check the detailed status information.
ℹ	The cluster is currently being initialized. Check the detailed status information.
⛔	Status call timed out. Server may be down. ▪ This message may be displayed if the node cannot be contacted. There may be a network issue affecting your management workstation, or the node may be offline. ▪ Refresh your web browser to check the connection to the node. ▪ If the problem persists, check the cluster status on the other node.

Technical considerations

Deploying an SSL certificate

Special consideration needs to be given to deployments that require SSL access to the cluster.

The Common Name (CN) of an SSL certificate must match the hostname of the site being visited. Certificates that do not meet this requirement may still be used to secure the connection, but a browser security warning is displayed. In modern browsers this warning is intended to deter users from what may be a potentially serious “man in the middle” attack. Non-technical visitors should not be expected to analyze and interpret these messages.

Where SSL access is a requirement, the recommended approach is to issue the certificate for the hostname of the cluster’s virtual IP address, and install the same certificate on both nodes.

This approach ensures that all operator and visitor access to the cluster is secured with a certificate that matches the hostname and IP address, avoiding any unnecessary browser security warnings.

NOTE When using this approach, the administrator will receive browser security warnings about the certificate hostname mismatch if he accesses each node individually.

3 Configuring High Availability

Check plugin versions

Deploying a High Availability cluster requires the following plugin versions:

- High Availability Services 0.9.14 or later

To verify you have the correct plugin versions installed, navigate to **Administrator** > **Plugin Manager** > **List Available Plugins** and check the version number in the list.

Use the **Check for Plugin Updates** link to download and install updated plugins.

Cluster setup

Before you begin, review this checklist to ensure you are prepared to setup a cluster:

- You have two servers available;
- Each server is powered up and connected to the same local area network;
- Each server has a unique hostname;
- Each server has a valid subscription ID and has been updated using the Plugin Manager. Ensure that the High Availability Services plugin has been installed along with any available plugin updates;
- You are logged in as the administrator on each server;
- You have determined the desired network configuration (virtual IP address) for the cluster.

Click the **Create New Cluster** command link to begin the process of creating a new cluster.

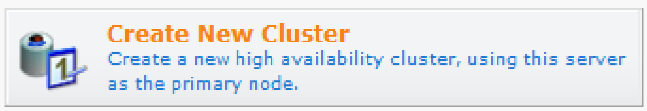

Create New Cluster

Create a new high availability cluster, using this server as the primary node.

Prepare primary node

Use the Cluster Configuration form to enter the basic network and control parameters for the cluster.

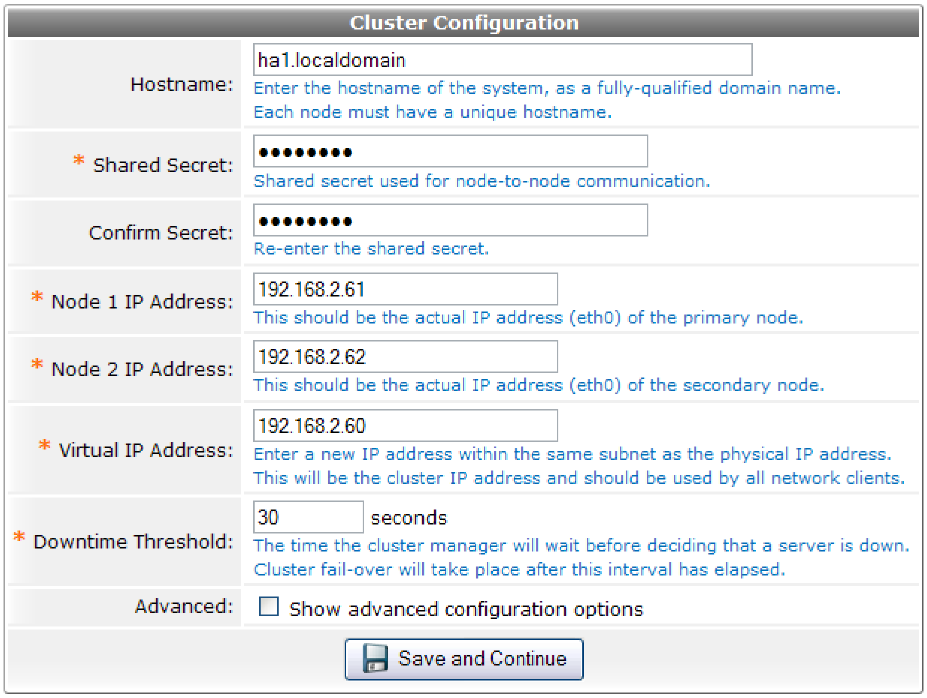

Cluster Configuration

Hostname:	ha1.localdomain Enter the hostname of the system, as a fully-qualified domain name. Each node must have a unique hostname.
* Shared Secret:	•••••••• Shared secret used for node-to-node communication.
Confirm Secret:	•••••••• Re-enter the shared secret.
* Node 1 IP Address:	192.168.2.61 This should be the actual IP address (eth0) of the primary node.
* Node 2 IP Address:	192.168.2.62 This should be the actual IP address (eth0) of the secondary node.
* Virtual IP Address:	192.168.2.60 Enter a new IP address within the same subnet as the physical IP address. This will be the cluster IP address and should be used by all network clients.
* Downtime Threshold:	30 seconds The time the cluster manager will wait before deciding that a server is down. Cluster fail-over will take place after this interval has elapsed.
Advanced:	☐ Show advanced configuration options

If you have not already set a unique hostname for this server, you can do so here. Each node in the cluster must have a unique hostname.

You must enter a shared secret for this cluster. The shared secret is used to authenticate the messages sent between the nodes in the cluster.

The downtime threshold parameter explained in section 0 may be specified on this form.

Click the **Save and Continue** button to prepare the primary node.

Prepare secondary node

To prepare the secondary node, log in to that node and click the **Join Cluster** command link.

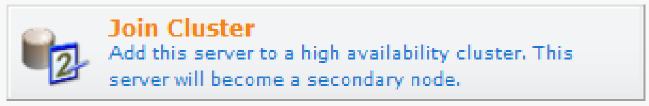

Join Cluster

Add this server to a high availability cluster. This server will become a secondary node.

Use the Cluster Configuration form to enter the shared secret for the cluster and the IP address of the primary node.

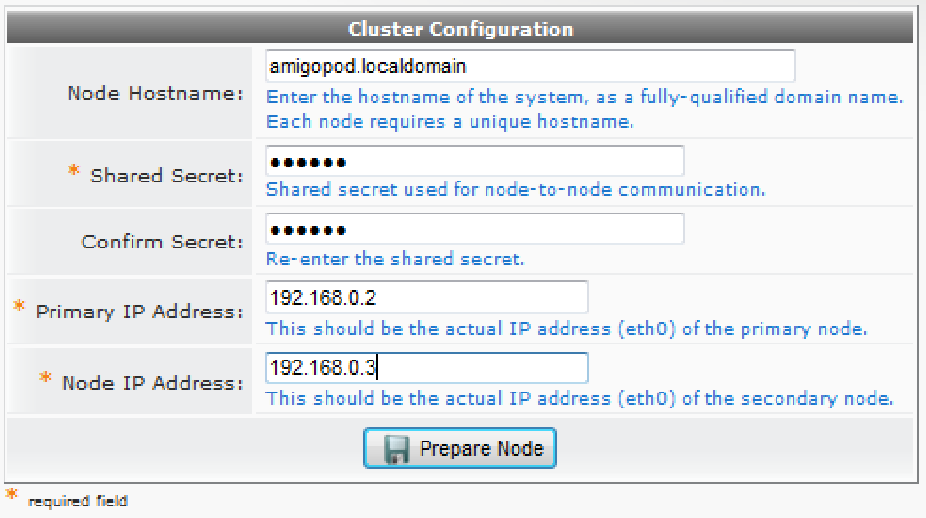

Cluster Configuration

Node Hostname:	amigopod.localdomain Enter the hostname of the system, as a fully-qualified domain name. Each node requires a unique hostname.
* Shared Secret:	●●●●●● Shared secret used for node-to-node communication.
Confirm Secret:	●●●●●● Re-enter the shared secret.
* Primary IP Address:	192.168.0.2 This should be the actual IP address (eth0) of the primary node.
* Node IP Address:	192.168.0.3 This should be the actual IP address (eth0) of the secondary node.

Click the **Prepare Node** button to save and verify the settings for the secondary node.

Cluster initialization

To complete the setup of the cluster, return to the primary node after preparing the secondary node and click the **Confirm Node Settings** button.

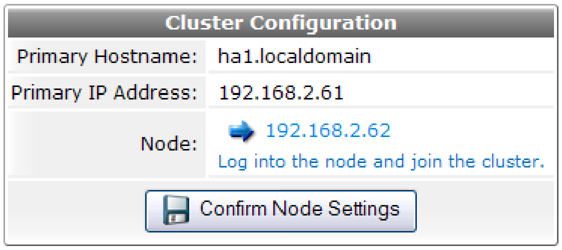

Cluster Configuration

Primary Hostname:	ha1.localdomain
Primary IP Address:	192.168.2.61
Node:	192.168.2.62 Log into the node and join the cluster.

The Cluster Initialization form is displayed.

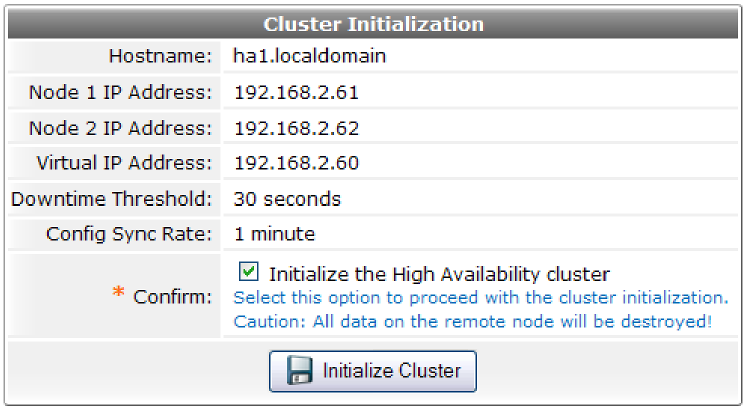

Cluster Initialization

Hostname:	ha1.localdomain
Node 1 IP Address:	192.168.2.61
Node 2 IP Address:	192.168.2.62
Virtual IP Address:	192.168.2.60
Downtime Threshold:	30 seconds
Config Sync Rate:	1 minute
* Confirm:	☑ Initialize the High Availability cluster Select this option to proceed with the cluster initialization. Caution: All data on the remote node will be destroyed!

Select the checkbox and click the **Initialize Cluster** button to proceed.

NOTE During the cluster initialization process, the entire contents of the RADIUS database (including guest accounts, user roles, and accounting history) and all configuration settings of the primary node will be replicated to the secondary node. The existing database contents and configuration settings on the secondary node will be destroyed. It is very important to ensure that you have selected the correct node as the primary node, particularly if you are rebuilding the cluster. If in doubt, it is recommended that you perform a complete backup of both nodes prior to initializing the cluster.

Several status messages and a progress meter will be displayed while the cluster is initialized, which may take several minutes depending on the amount of data to be replicated.

Once the initialization process completes, you will be returned to the High Availability start page, where the cluster status will be displayed as:

The cluster is running normally.

Cluster deployment

After setting up a cluster, you must make appropriate configuration changes for your network to take advantage of the cluster's fault tolerance.

The principal configuration change required is to replace the IP address of a single Amigopod server with the virtual IP address of the cluster.

- NAS devices and other RADIUS clients should be configured with the cluster IP address.
- Operators should use the cluster's IP address when provisioning guest accounts.
- Configure NAS devices to redirect visitors to the cluster's IP address for web login pages. Only the IP address in the redirection URL should be changed; the remainder of the redirection URL should not be altered.

The network administrator should use the node IP addresses to perform system administration tasks on each node, including managing the cluster itself.

Cluster maintenance

Use the **Cluster Maintenance** command link to access maintenance functions related to the cluster.

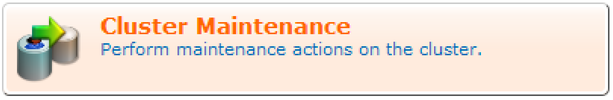

Cluster Maintenance
Perform maintenance actions on the cluster.

The maintenance commands that are available on this page will depend on the current state of the cluster as well as which node you are logged into.

NOTE Some maintenance commands are only available on the secondary node. Other commands may change the active state of the cluster. For this reason it is recommended that cluster maintenance should only be performed by logging into a specific node in the cluster using its IP address.

Recovering from a failure

From a cluster maintenance perspective, there are two kinds of failure:

- A **temporary outage** is an event or condition that causes the cluster to fail-over to the secondary node. Clearing the condition allows the cluster's primary node to resume operations in essentially the same state as before the outage.

- A **hardware failure** is a fault that to correct requires rebuilding or replacing one of the nodes of the cluster.

The table below lists some system failure modes and the corresponding cluster maintenance that is required.

Failure Mode	Maintenance
Software failure – system crash, reboot or hardware reset	Temporary outage
Power failure	Temporary outage
Network failure – cables or switching equipment	Temporary outage
Network failure – appliance network interface	Hardware failure
Hardware failure – other internal appliance hardware	Hardware failure
Data loss or corruption	Hardware failure

Recovering from a temporary outage

Use this procedure to repair the cluster and return to a normal operating state:

1. This procedure assumes that the primary node has experienced a temporary outage, and the cluster has failed over to the secondary node.

2. Ensure that the primary node and the secondary node are both online.

3. Log into the secondary node. (Due to fail-over, this node will be assigned the cluster's virtual IP address.)

4. Click Cluster Maintenance, and then click the **Recover Cluster** command link.

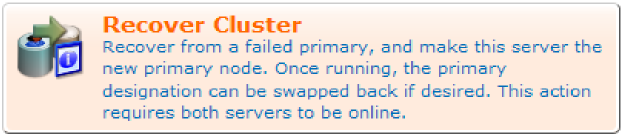

Recover Cluster

Recover from a failed primary, and make this server the new primary node. Once running, the primary designation can be swapped back if desired. This action requires both servers to be online.

5. A progress meter is displayed while the cluster is recovered.

NOTE The cluster's virtual IP address will be temporarily unavailable while the recovery takes place.

6. Recovery is complete. The secondary node is now the new primary node for the cluster. The cluster is back in a fault-tolerant mode of operation.

The Recover Cluster command will only work if the node that failed is brought back online with the same cluster configuration. This is normally the case in all temporary outages. If this is not the case, use the procedure described in section 0 to recover the cluster.

To return the primary node back to its original status as the primary node in the cluster, you can use the Swap Primary Servers command described in section 0.

Recovering from a hardware failure

If the failed node has been replaced, the cluster configuration will no longer be present on that node. To recover the cluster, first ensure that the replaced node is ready to rejoin the cluster, then destroy the cluster and recreate it.

Use the following procedure to rebuild the cluster:

This procedure assumes that the primary node has failed and has been replaced.

1. Configure the network settings, subscription IDs and hostname for the replacement primary node.

2. Ensure that the replacement primary node and the secondary node are both online.

3. Log into the secondary node. (Due to fail-over, this node will be assigned the cluster's virtual IP address.)

4. Click Cluster Maintenance, and then click the **Destroy Cluster** command link.

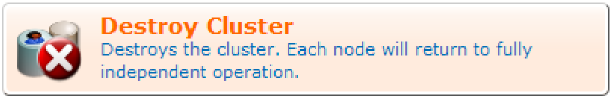

Destroy Cluster
Destroys the cluster. Each node will return to fully independent operation.

5. A progress meter is displayed while the cluster is destroyed.

NOTE The virtual IP address of the cluster will be unavailable until the cluster is reinitialized.

6. Click the **Create New Cluster** command link.

7. Recreate the cluster using the process described in section 3. Note that the new cluster's primary node must be the former cluster's secondary node that you are presently logged into.

8. When the cluster is initialized, the database and configuration is replicated to the replacement primary node.

9. Recovery is complete. The cluster's virtual IP address is now available, and the secondary node is now the new primary node for the cluster. The cluster is back in a fault-tolerant mode of operation.

A similar procedure can be used to rebuild the cluster in the event of a secondary node suffering a hardware failure.

Performing scheduled maintenance

Routine maintenance tasks such as a server reboot or shutdown may occasionally be required for a server that is part of a cluster.

These tasks may be performed by ensuring that the server is the secondary node in the cluster. If the secondary node goes offline, the primary node will be unaffected and the cluster will continue to provide network services without interruption. When the secondary node comes back online, the cluster will be automatically rebuilt and replication will resume.

To check the current status of a node, log into that node and click the ℹ **Show details** link displayed with the cluster status on the High Availability page. The node's current status is displayed under the **Local Status** heading.

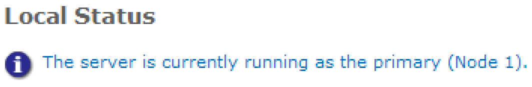

Local Status

ℹ The server is currently running as the primary (Node 1).

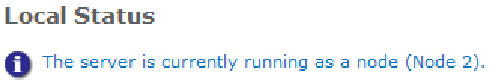

Local Status

ℹ The server is currently running as a node (Node 2).

Use this procedure to make the current primary node the secondary node:

1. Log into the current secondary node of the cluster.

2. Click Cluster Maintenance, and then click the **Swap Primary Server** command link.

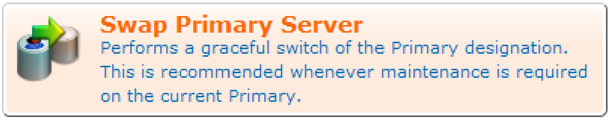

Swap Primary Server
Performs a graceful switch of the Primary designation. This is recommended whenever maintenance is required on the current Primary.

3. A progress meter is displayed while the primary node is switched.

NOTE The cluster's virtual IP address will be temporarily unavailable while the swap takes place.

4. The swap is complete. The secondary node is now the new primary node for the cluster. The cluster is back in a fault-tolerant mode of operation.

5. Perform any required maintenance on the new secondary node.

Destroying a cluster

The **Destroy Cluster** command link is used to shut down a cluster and return to independent nodes.

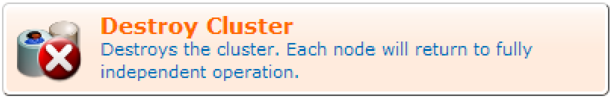

Destroy Cluster
Destroys the cluster. Each node will return to fully independent operation.

Immediately after the cluster is destroyed, both nodes will have the same database and configuration state. However, changes on one node will no longer be replicated to the other node as the cluster is no longer functioning.

NOTE Avoid using this command when you are accessing the cluster using its virtual IP address, as the virtual IP address will no longer be available when the cluster has been destroyed.

Cluster troubleshooting

When building a cluster, use the recommended values for the downtime threshold, keep-alive rate and configuration sync rate. You should only change these values if you have a specific requirement and have verified that different values can be used to meet that requirement.

To avoid unexpected fail-over of the cluster, ensure that the network connection to the nodes of the cluster is always available. Use high quality network equipment, including cables, and secure physical access to the servers to prevent accidental dislodgement of cables.

If network access to the cluster is intermittent, this may indicate a possible hardware failure on the current primary node. In this situation, you may either use the Swap Primary Server command to make the secondary node the new primary node, or you can cause the cluster to fail-over to the secondary by disconnecting the primary node.

Brief network outages are permissible and will not cause fail-over, provided that the network outage is shorter than the downtime threshold of the cluster.

During a fail-over from the primary to the secondary node, the network services provided by the cluster will be unavailable. The time that the cluster will be offline is bounded by the downtime threshold. This can be used to calculate the expected availability of the cluster.

The **Restart Cluster Services** and **Stop Cluster Services** command links on the Cluster Maintenance page may be used to test fail-over conditions by simulating a cluster failure.

NOTE Avoid using these commands when you are accessing the cluster using its virtual IP address, as the virtual IP address may become unavailable.

The **View Log Files** command link allows the internal state of the cluster to be viewed.

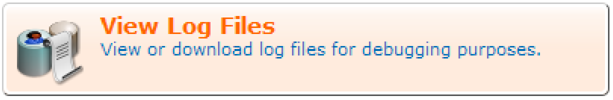

View Log Files
View or download log files for debugging purposes.

This may be useful if debugging a problem related to the cluster. The log files may be exported to a zip file. If you require support about a cluster-related problem, include a copy of the exported cluster log files with your support request.