Clustering Configuration Guide

Published August 2020

Document Version 1.0

Pulse Secure, LLC
2700 Zanker Road,
Suite 200 San Jose
CA 95134

www.pulsesecure.net

Clustering Configuration Guide

The information in this document is current as of the date on the title page.

END USER LICENSE AGREEMENT

The Pulse Secure product that is the subject of this technical documentation consists of (or is intended for use with) Pulse Secure software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at http://www.pulsesecure.net/support/eula/. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Contents

Clustering

Clustering Feature Overview

Clusters define a collection of servers that operate as if they were a single machine. A cluster pair is used to refer to a cluster of two units and a multiunit cluster refers to a cluster of more than two units. Once two or more units are joined in a cluster, they act as one unit.

Figure 1 shows two PSA series devices deployed as a cluster pair.

Figure 1 Clustering

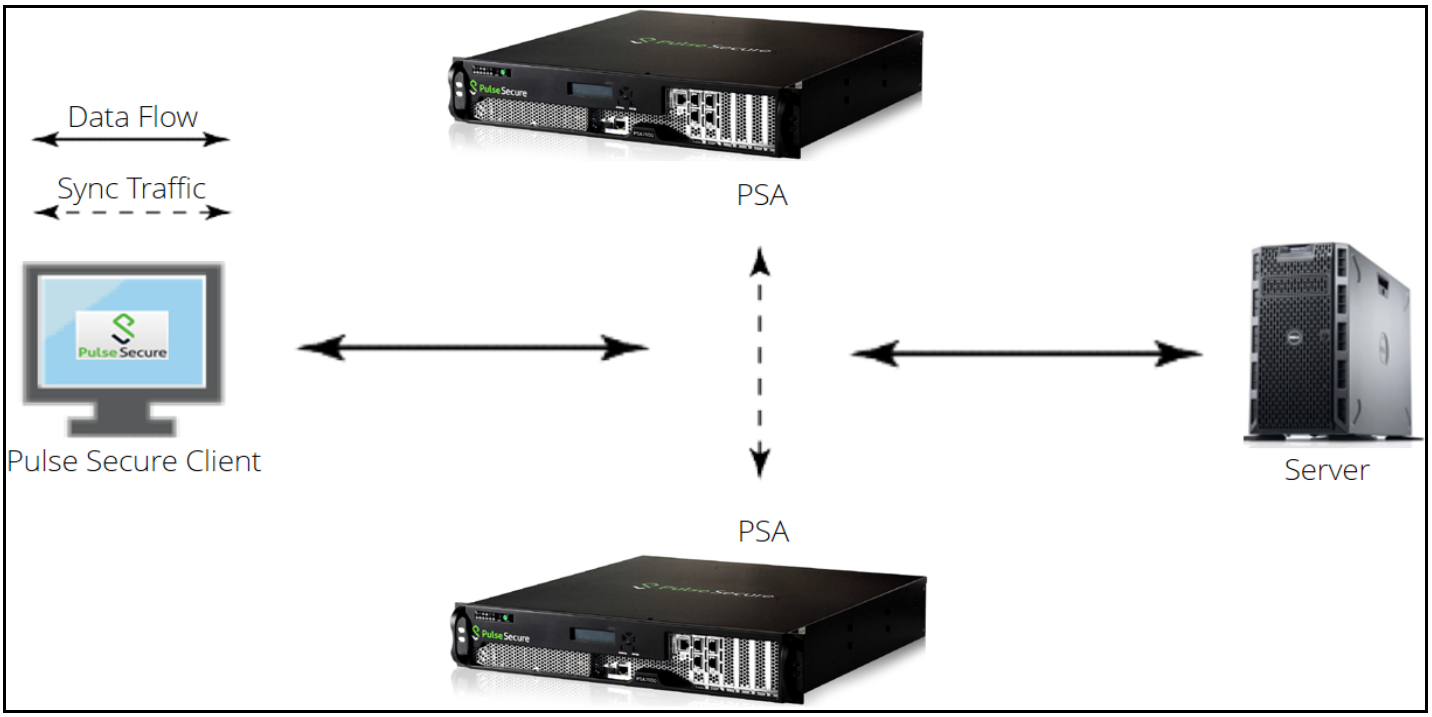

Deployments

Pulse Secure access management framework supports two types of clusters:

- Load balancing clusters or active/active clusters
- Failover clusters or active/passive clusters

Load balancing clusters or active/active clusters - Load balancing clusters provide scalability and increase availability of Web-based services. Figure 2 shows an example of an active/active deployment. A user can deploy 4 node cluster on PSA-7000. All other platform models support 2 node clusters only.

Note: The system (UI) allows adding up to 8 nodes. However, only up to 4 nodes in a cluster have been officially qualified.

Failover clusters or active/passive clusters - Failover clusters provide high availability (HA). The primary purpose of HA clusters is to provide uninterrupted access to data, even if a server loses network or storage connectivity, or fails completely, or if the application running on the server fails. Figure 2 shows an example of an active/passive deployment. The active/passive cluster supports only 2 node clusters in all types of platforms except VA.

Note: For further information on clustering and scalability, please contact Pulse Secure technical help.

Note: Pulse Secure access management framework also supports an IPv6 configuration for active/active and active/passive clusters.

Figure 2 Active/Active Deployment

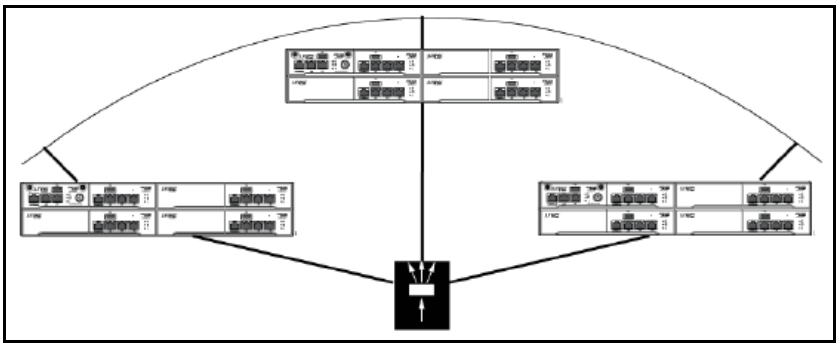

Figure 3 Active/Passive Deployment

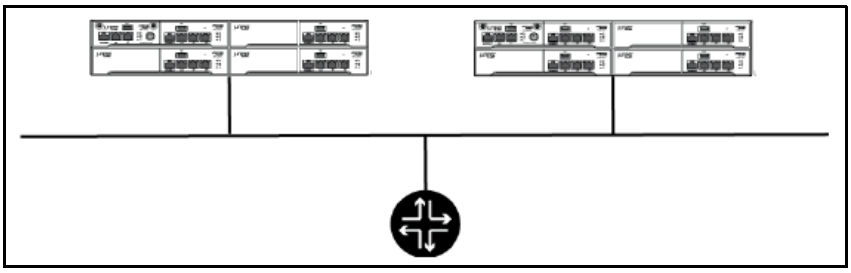

Requirements and Limitations

You must follow these considerations when deploying a cluster:

- Cluster members must run the same software version.
- Cluster members must use the same hardware platform.
- State synchronization must occur only through the internal Network Interface Card (NIC).
- Ensure the cluster communication and resource access must take place over an internal network.
- You can deploy an active/passive clustering only within the same IP subnet.

Cluster Licensing

Pulse Connect Secure devices share licenses within the cluster.

Administrator can:

- create license server with active/active cluster on virtual/cloud and hardware platforms
- lease all different types of licenses to license clients from any node of the active/active cluster.
- surrender/recall licenses from any node of the active/active cluster

Key points about licenses in a cluster:

Within a cluster, user licenses are shared among all nodes.

Licenses are additive in cluster. The total user count would be sum of all nodes in the cluster.

Install licenses equally on each node in the cluster

Careful consideration should be taken when removing a node from a cluster. Removing a node will impact the total number of users.

Note: The nodes in a virtual appliance cluster needs to have the same virtual appliance core licenses.

Reason for installing licenses equally in a cluster

The reason is to prevent loss of user count during node failure. A node can only borrow up to two times (2x) the total number of licenses installed on the device locally for a 10-day grace period. After 10-day period, the user count will revert to the total number of licenses installed locally.

If the licenses are evenly distributed in the cluster, the user count will remain the same regardless which node fails. If the license is unevenly distributed in the cluster, the user count will be different depending on which node fails (Refer to Scenario A and B for examples).

For example:

Node_108 = 35 user license

Node_109 = 100 user license

When clustered, the total number of users is 135 (35 + 100 = 135).

Scenario A:

If Node_109 goes down, the total number of licenses would be 70 users. (35 X 2 = 70).

The maximum number of users will be 70 for the 10-day grace period. After the grace period expires, the user count will drop to 35 users.

Scenario B:

If Node_108 goes down, the total number of licenses would be 135 users (100 + 35 = 135).

The maximum number of users will be 135 users for the 10-day grace period. After the grace period expires, the user count will drop to 100 users.

Why is the calculation different?

In this scenario, the device can only borrow up to the total number of user licenses in the cluster. This means instead of (100 x 2 = 200), it will only (100 + 35 = 135).

How License count is impacted when removing a node from a cluster?

If a node is removed from an existing cluster, this will decrease the total user count. For example:

Node A = 100

Node B = 100

Node C = 100

The total user count would be 300 (100 + 100 + 100). If Node C is removed from the cluster, this will drop the total number to 200 (100 + 100) as Node C is no longer part of the cluster.

Recommendation for node replacement:

The general recommendation is to never delete a node from a cluster unless the device needs to be changed to a standalone device. If a node needs to be replaced in the existing cluster, the following steps are recommended:

Note: The following steps should be performed during a maintenance window. When the replacement device joins the cluster, this will cause the web server to restart causing a short disruption to connected users.

1. Power on the replacement device and complete the initial configuration
2. Log in to the admin console
3. Importing the existing system and user configuration to the replacement device
4. Install new licenses on the replacement device. If this is an RMA device, complete the RMA process to install replacement licenses on the device.

During the import of the system and user configuration, it will retain the cluster configuration and allow the replacement device to join the existing cluster.

Deploying an Active/Active Cluster

This example describes the tasks involved in deploying an active/active cluster. It includes the following information:

Overview

An active/active clustering provides high availability and load balancing when deployed with an external load balancer. An active/active cluster deployment requires an external device to distribute the load among the members because the cluster does not have a VIP address. The load balancing devices are equipped with algorithms that balance the load, as well as detect whether a device is down.

Active/active configuration allows increased aggregate system throughput as well as seamless failover, which is achieved by state synchronization between the two devices for all the configurations so that the devices are virtually identical. Figure 4 shows active/active clustering deployed with an external load balancer.

Note: This feature provides increased throughput and performance for peak load characteristics; however, it does not provide increased scalability beyond the total licensed users.

Network Topology

Active/active clustering can support up to eight nodes in a cluster but are also supported in a LAN environment. Within an active/active cluster, no VIP address is present, and each cluster member has its own network settings. Figure 4 shows an example of active/active deployment.

Figure 4 Active/Active Clustering

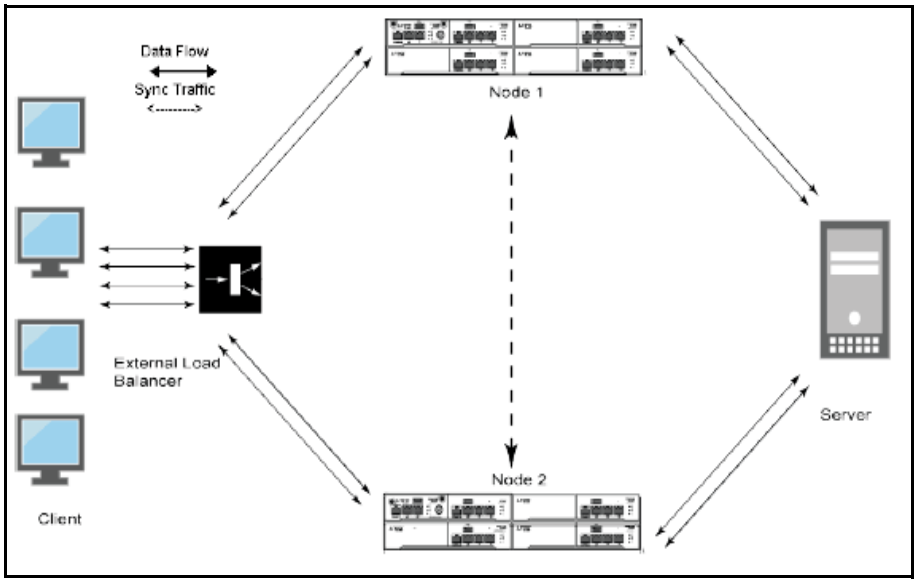

Before You Begin

We recommend that you deploy a cluster in a staging environment first and then move to a production environment after testing the authentication realm, user role, and resource policy configurations, as well as any applications your end users might access.

You must follow these considerations when deploying a cluster:

- Cluster members must run the same software version.
- Cluster members must use the same hardware platform.
- State synchronization must occur only through the internal Network Interface Card (NIC).

- Ensure the cluster communication and resource access must take place over an internal network.
- For better performance, consider the following recommendations:
  - Use Dual arm configuration.
  - Do not use source NAT configuration.
  - Disable Multicast unless necessary.

When choosing and configuring a load balancer for your cluster, we recommend that you ensure the load balancer:

- Supports IPsec
- Listens for traffic on multiple ports
- Can be configured to manage traffic using assigned source and destination IP addresses (not destination port)

Configuring an Active/Active Cluster

You use the primary node admin GUI to create the cluster and add members. The primary node is added as part of the cluster creation operation. When you add members, you are prompted for settings unique to the member, such as the name and IP address configuration for the internal and external interfaces. A few additional settings are also unique, namely the management port and VLAN port settings, so you add these manually after the add node procedure that follows, but before the join cluster operation.

Note: If IPv6 is required, then configure both the nodes with IPv6 settings before creating the cluster.

To create a cluster and add members:

1. Select **System > Clustering > Create Cluster** and enter a name for the cluster, a cluster password, and a name for this node, such as Node-X.

   You need to enter the password again when specifying additional nodes to join the cluster. All nodes in the cluster use this password to communicate.

Figure 5 shows the Create New Cluster page for Pulse Connect Secure.

Figure 5 Create New Cluster Page

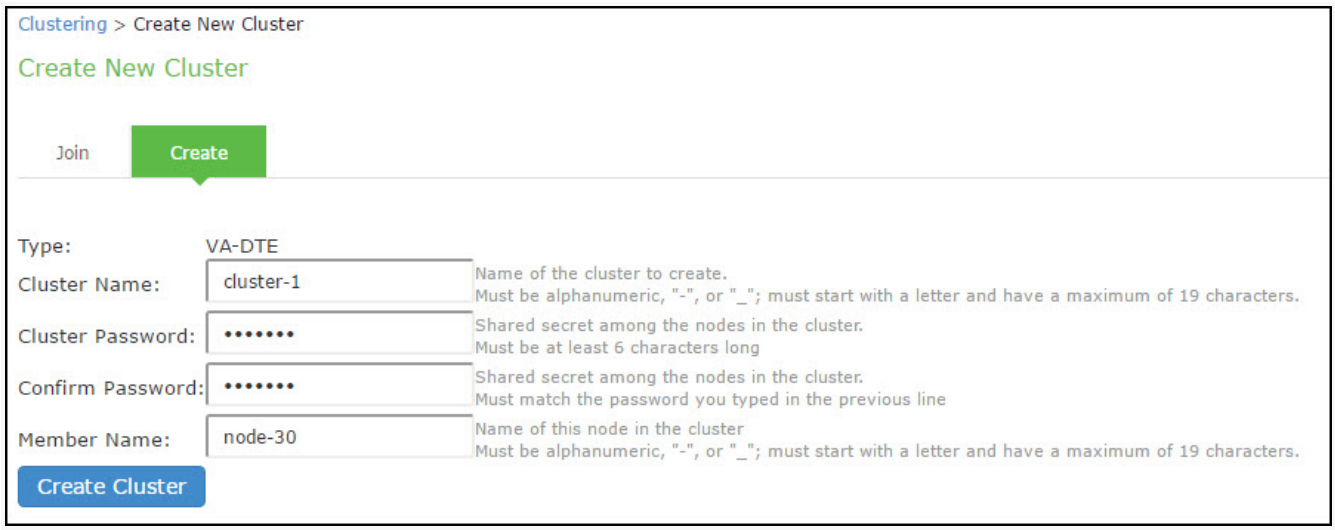

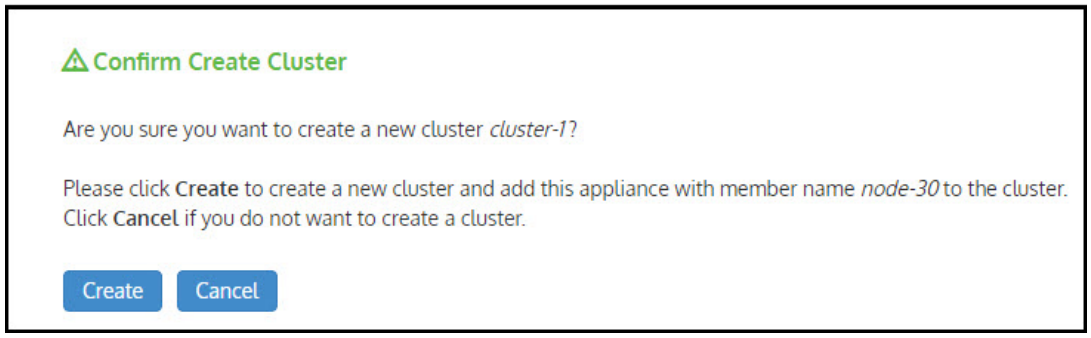

2. Click **Create Cluster.** When prompted to confirm the cluster creation, click **Create**. After the device initializes the cluster, the Clustering page displays the **Status** and **Properties** tabs.

3. Click **Properties**.

Figure 6 shows the Clustering page for Pulse Connect Secure.

Figure 6 Clustering Page- Active/Active Configuration

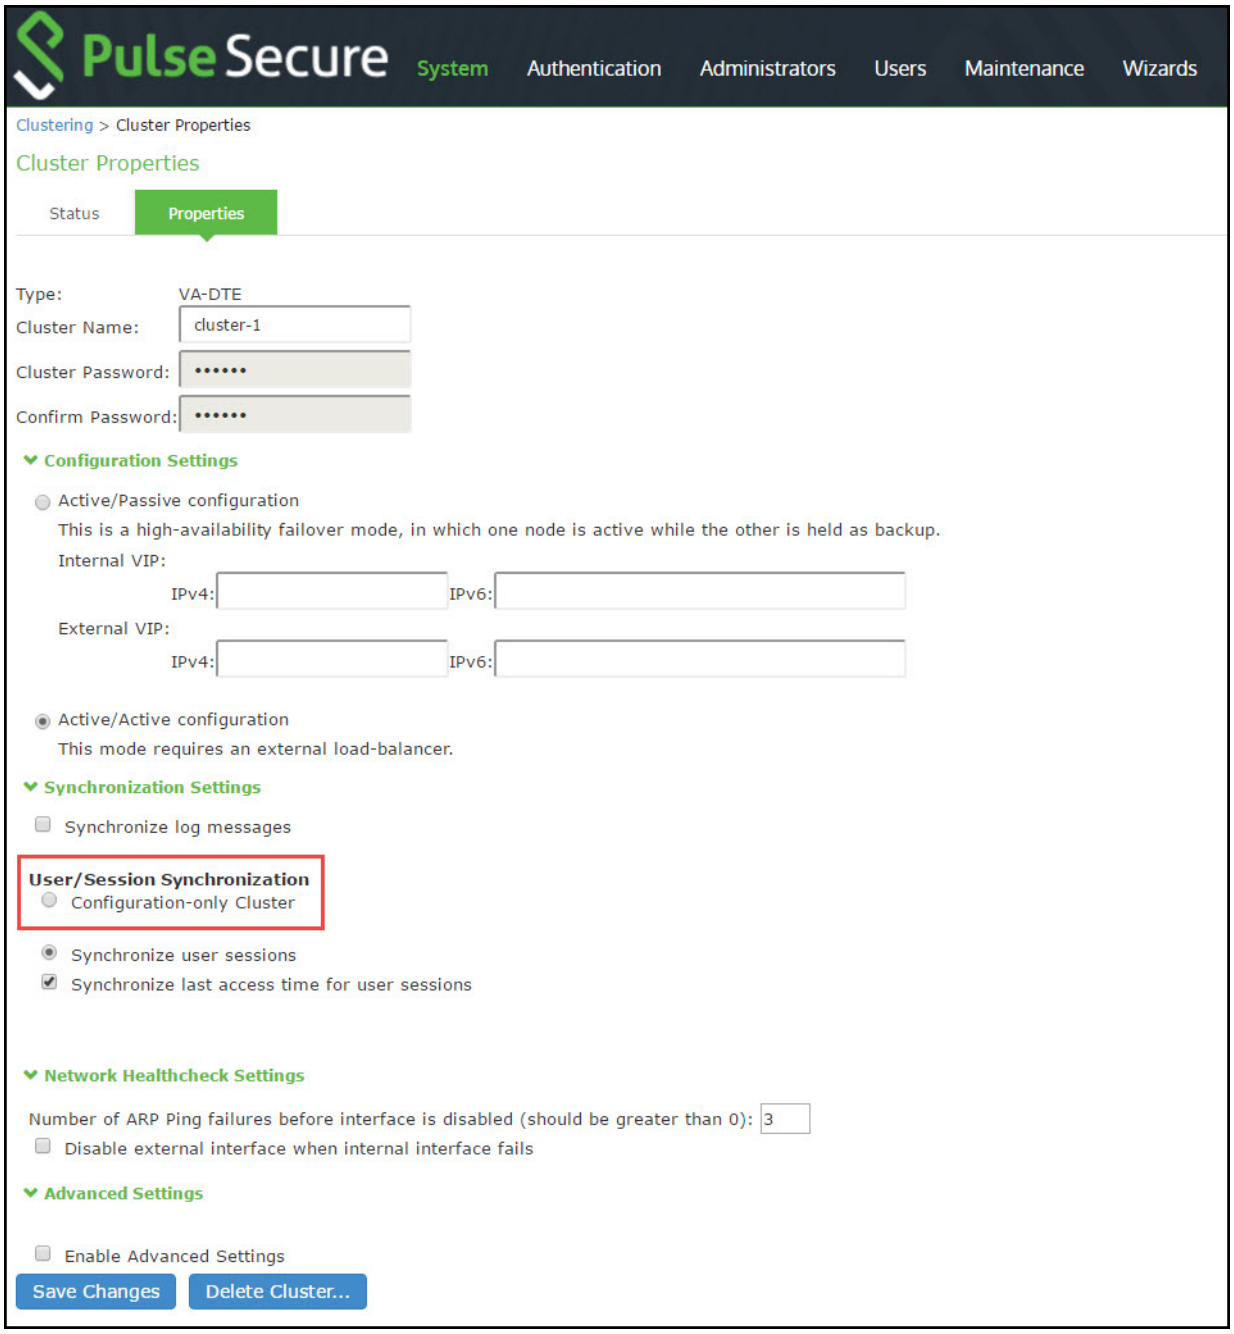

4. Select **Active/Active configuration** and complete the configuration as described in Table 1 Active/Active configuration is selected by default.

Table 1 Clustering Property Settings

Settings	Guidelines
Cluster Name	Specifies a name to identify the cluster.
Configuration Settings	
Active/Passive configuration	Select this option to run a cluster pair in active/passive mode. Then, specify an internal VIP (virtual IP address) and an external VIP if the external port is enabled.
Active/Active configuration	(Default) Select this option to run a cluster pair in active/active mode. Active/Active runs a cluster of two or more nodes in active/active mode using an external load balancer. **Note:** To change a two-unit active/passive cluster to an active/active cluster with more than two nodes, first change the configuration of the two-unit cluster to active/active and then add the additional nodes.
Synchronization Settings	
Synchronize log messages	Select this option to propagate all log messages among the devices in the cluster.
User/Session Synchronization	
Configuration only cluster	Select this option to disable synchronization of session data and to replicate only configuration data and user records (for example, web bookmarks, NFS and windows shared files, terminal servers, telnet sessions, SAM, preferences, and passwords). **Note:** Enabling this option limits data transfers between the cluster nodes. User and Session specific limits are only enforced on the node and not across the cluster.
Synchronize user sessions	Select this option to synchronize all user session information (for example, instances of access to intranet services) among all the devices in the cluster.
Synchronize last access time for user sessions	Select this option to propagate the latest user access information across the cluster.
Note: If you select both Synchronize log messages and Synchronize user sessions check boxes, everything is replicated on the cluster nodes, including networking information. Even though networking information, including syslog and SNMP settings, can be configured per node or per cluster, all of the networking information is synchronized between nodes when these two options are set. If your cluster node configurations diverge because of changes made to one node while another is disabled or unavailable, the system manages the remerging of the configurations automatically, for up to 16 updates. Beyond the maximum number of allowable updates, you might need to intervene and remerge the configurations manually. In some instances, the system might be unable to remerge the configurations if there is not enough overlapping configuration information between two nodes to manage the internode communication. For example, for a two-node cluster in which the two nodes are partitioned from each other because of a network outage, if the internal network IP address of one of the nodes changes in one of the partitions, the two partitions are unable to rejoin, even when the network is repaired. In such a case, you must remerge the configurations manually.	
Network Healthcheck Settings	
Number of ARP Ping Failures	Specify the number of ARP ping failures allowed before the internal interface is disabled.
Disable external interface when internal interface fails	Select this option to disable the external interface of the device if the internal interface fails.

Settings	Guidelines
Advanced Settings	
Enable Advanced Settings	Select the **Advanced Settings** check box to specify the timeouts for the underlying cluster system. Do not change any values under this setting unless instructed to do so by Pulse Secure Technical Support.
Network Type	Select the appropriate network type. Network type selection controls the timeouts used by the underlying cluster system. Change this value only when you observe repeated cluster partitions that may be related to long network delays or significant load in any of the cluster nodes. A non-default network type cannot be used in conjunction with non-default timeout multipliers. If a non-default network type is picked, the timeout multiplier will silently get reset to the default value.
Timeout Multiplier	Default cluster timeouts have been picked to be optimal for typical cluster installations. Administrators have the ability to adjust the cluster timeouts over a linear scale of 1-20. Smaller timeouts result in faster failure detection. Larger timeouts minimize the risk of cluster splits during transient network glitches. The system can be instructed to pick a reasonable default for the current cluster configuration by specifying a value of 0. A non-default timeout multiplier can only be used in conjunction with the default network type. If a non-default network type is picked, the timeout multiplier will silently get reset to the default value.

5. Click **Save Changes.**
6. Click **Add Members** to specify additional cluster nodes.
7. Click **Save Changes.**
8. Select **System > Network > Management Port > Settings** and configure the management port IPv4 and IPv6 (if configured) of node-2.

   Figure 7 Configuring Management Port

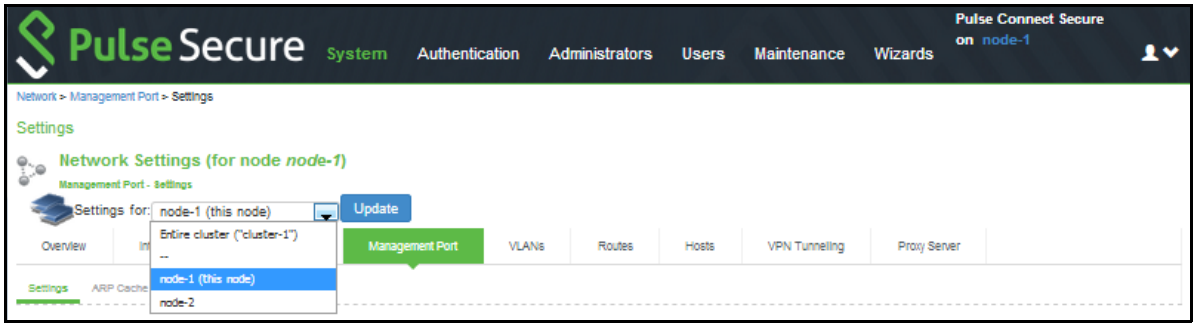

9. If a license server needs to be configured on both the nodes of a cluster, then perform the following steps:
   a. Navigate to **Configuration > Licensing > Configure Server.**
   b. Select the setting for **Entire cluster.**
   c. Configure the License server IP and preferred network.

d. Click **Save Changes.**

Figure 8 Configuring License Server for Entire Cluster

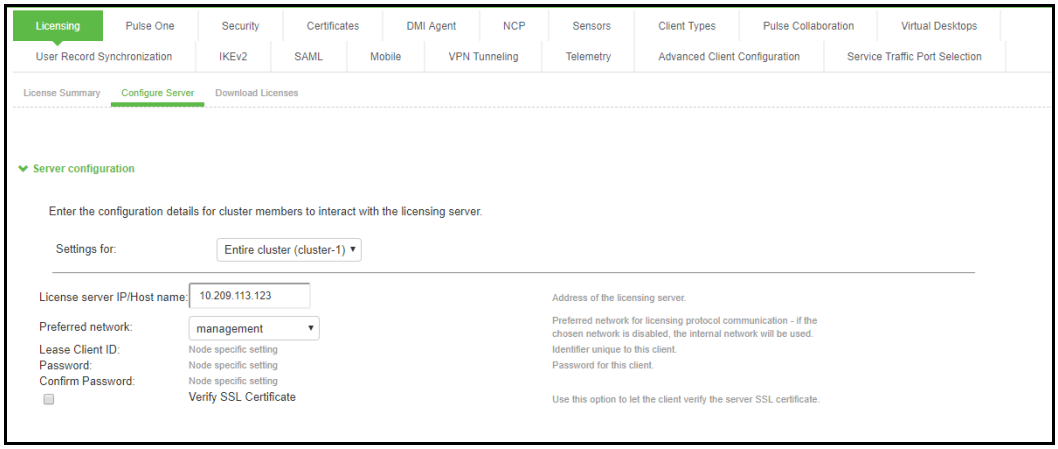

e. Now, select the settings for node-wise and provide **Lease Client ID, Password and Çonfirm Password for each node.**

Figure 9 Node-wise Server Configuration

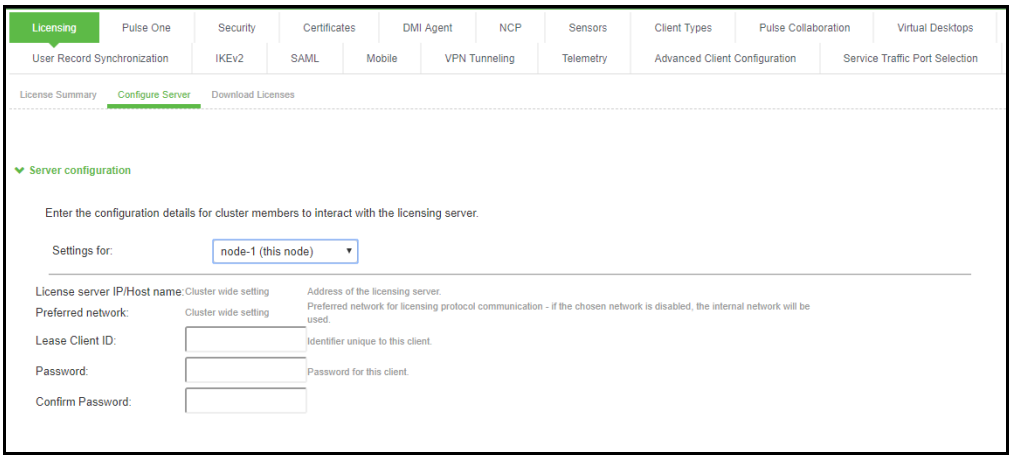

Joining Nodes to the Cluster

The primary node joins the cluster as part of the creation process.

To join additional nodes to the cluster:

1. From an existing cluster member, select the **System > Clustering > Cluster Status** tab and specify the node you want to add to the cluster.
2. From the admin console of the node you want to add to a cluster:
   a. Select the **System > Clustering > Join** tab and enter:
      - The name of the cluster to join
      - The cluster password you specified when defining the cluster
      - The IP address of an active cluster member

Figure 10 shows the configuration page for Pulse Connect Secure.

Figure 10 Join Existing Cluster

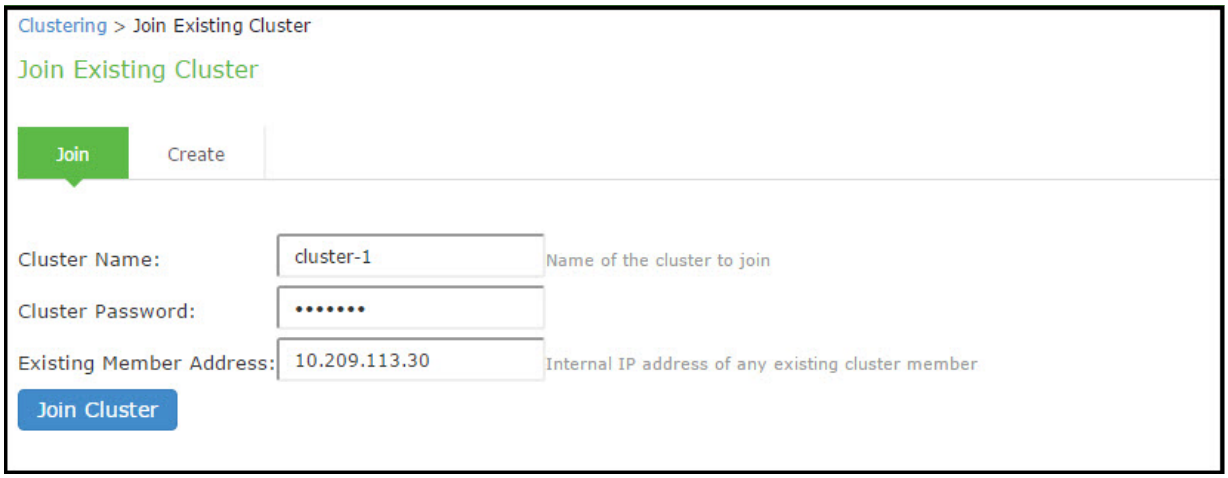

3. Click **Join Cluster**. When prompted to confirm joining the cluster, click **Join**.

While the new node synchronizes its state with the existing cluster member, each node's status indicates **Enabled**, **Enabled**, **Transitioning**, or **Enabled**, **Unreachable**.

When the node finishes joining the cluster, its Clustering page shows the Status and Properties tabs. After the node joins the cluster, you might need to sign in again.

Verifying

Purpose	Verifying the configuration on **System > Clustering > Cluster Status page.**
Action	Select **System > Clustering > Cluster Status.**

Figure 11 shows the status on the Clustering page for Pulse Connect Secure.

Figure 11 Clustering Page - Status

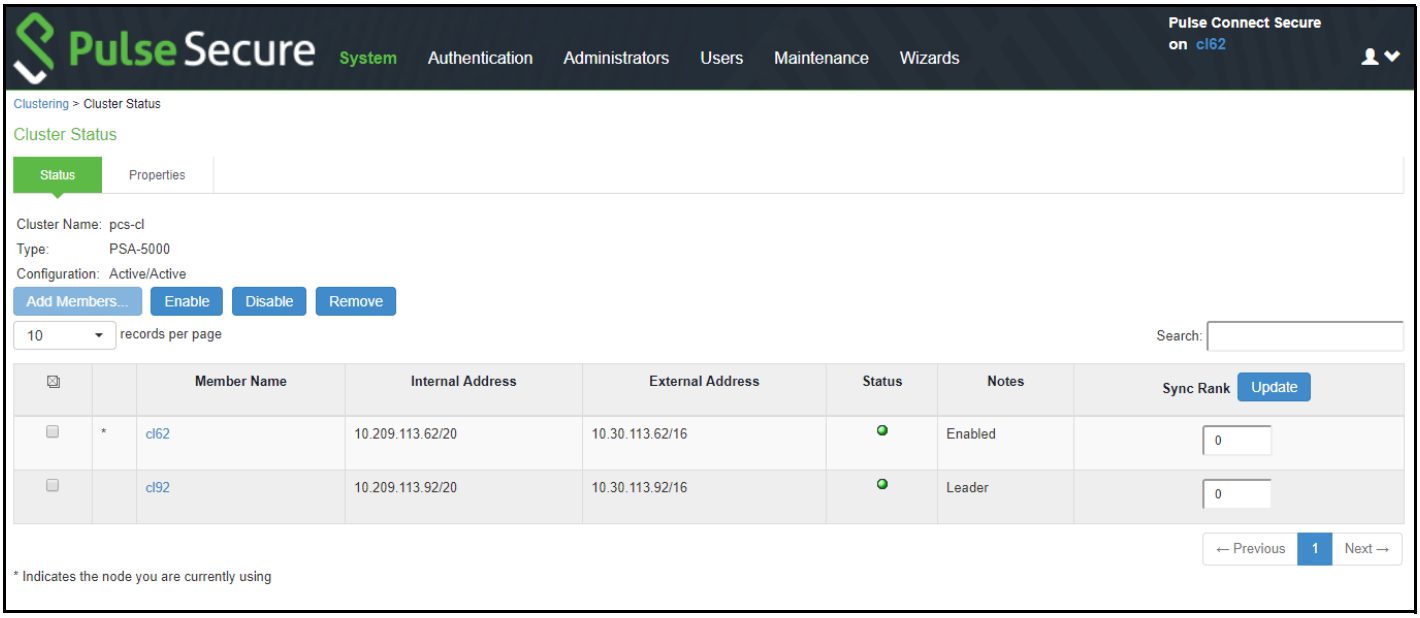

Table 2 describes the information displayed on the **Status** tab and the various management tasks you can perform, such as disabling, enabling, and removing a node from a cluster.

Table 2 Clustering Status

GUI Element	Description
Status Information labels	Displays the cluster name, type, configuration, internal VIP, and external VIP for an active/passive cluster only and not applicable for active/active cluster.
Add Members button	Click this button to specify a node you intend to add to the cluster. You can add multiple nodes at the same time.
Enable button	Click this button to enable a node that was previously disabled. When you enable a node, all state information is synchronized on the node.
Disable button	Click this button to disable a node within the cluster. The node retains awareness of the cluster but does not participate in state synchronizations or receive user requests unless members sign in to the node, directly.
Remove button	Click this button to remove the selected node or nodes from the cluster. After removal, the node runs in standalone mode.
Member Name column	Lists all nodes belonging to the cluster. You can click on a node's name to modify its name and network settings.
Internal Address column	Shows the internal IP address of the cluster member using Classless Interdomain Routing (CIDR) notation.
External Address column	Shows the external IP address of the cluster member using CIDR notation. Note that this column shows only the external IP address of the cluster leader unless you specify a different address for the node on its individual network settings page, which is accessible by clicking its name in the Member Name column. If you change the external IP address on the Network > Network Settings page, the change affects all cluster nodes.
Status column	Shows the current state of the node: • **Green light, Leader** - The node is the active member of an active/active cluster and is handling user requests. • **Green light, Enabled** - The node is handling user requests and participating in cluster synchronization. • **Yellow light, Transitioning** - The node is joining the cluster. • **Red light, Disabled** - The node is not handling user requests or participating in cluster synchronization. • **Red light, Enabled, Unreachable** - The node is enabled but because of a network issue, it cannot be reached. **Note:** A node's state is considered standalone when it is deployed outside of a cluster or after being removed from a cluster.

GUI Element	Description
Notes column	Shows the status of the node's connection to the cluster: • **OK** - The node is actively participating in the cluster. • **Transitioning** - The node is switching from the standalone state to the enabled state. • **Unreachable** - The node is not aware of the cluster. A cluster member might be unreachable even when it is online and can be pinged. Possible reasons include: • Incorrect password. • No information about all cluster nodes. • Configured with a different group communication mode. • Running a different service package version, or the machine is turned off.
Sync Rank column	Specifies the synchronization order for nodes when a node rejoins a cluster. Accepts sync ranks from 0 (lowest rank) to 255 (highest rank). The highest rank takes precedence. If two nodes have identical sync ranks, the alphanumeric rank of the member name is used to determine precedence.
Update button	Updates the sync rank after you change the precedence of the nodes in the Sync Rank column

Deploying an Active/Passive Cluster

This example describes the tasks involved in deploying an active/passive cluster. It includes the following information:

Overview

Active/passive clustering is supported only if the members of the cluster pair are in the same subnet because the VIP address must be shared by both the members. An active/passive cluster configuration provides high availability. Active/passive configurations allows seamless failover without the need to set up any external equipment, which is achieved by state synchronization between the two devices for all the configurations so that the devices are virtually identical. The Pulse Secure access control service uses a virtual IP (VIP) address to address the cluster pair in addition to addressing each device. The IP address takeover (IPAT) approach is used for the VIP address. If the active node fails, the passive node takes over the VIP address and sends a gratuitous Address Resolution Protocol (ARP) message notifying other networking devices that it now owns the VIP address. You should check that other devices in your network, especially the next-hop gateways, will honor the gratuitous ARP messages.

Topology

Figure 12shows active/passive clustering.

This feature provides increased throughput or capacity but does create redundancy in the case of a failure.

Figure 12 Active/Passive Clustering

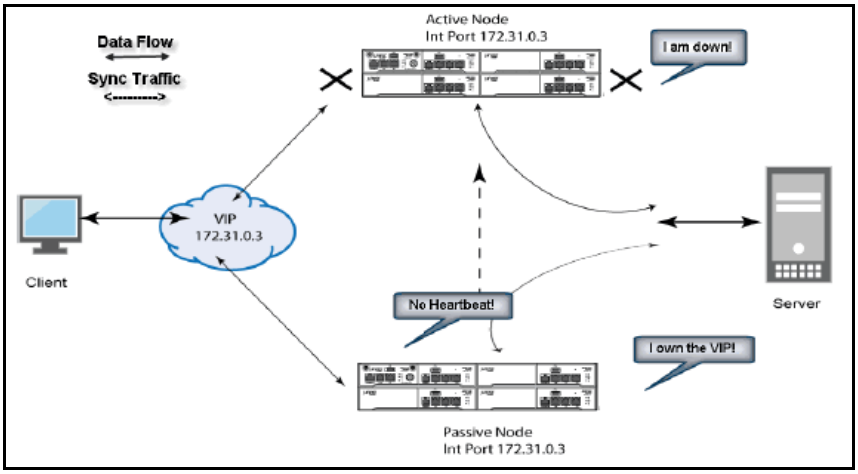

Requirements

Before you begin:

We recommend that you deploy a cluster in a staging environment first and then move to a production environment after testing the authentication realm, user role, and resource policy configurations, as well as any applications your end users might access.

You must follow these considerations when deploying a cluster:

- Cluster members must run the same software version.
- Cluster members must use the same hardware platform.
- State synchronization must occur only through the internal Network Interface Card (NIC).
- Ensure the cluster communication and resource access must take place over an internal network.

When choosing and configuring a load balancer for your cluster, we recommend that you ensure the load balancer:

- Supports IPsec
- Listens for traffic on multiple ports
- Can be configured to manage traffic using assigned source and destination IP addresses (not destination port)

Guidelines and Limitations

- A virtual IP address (VIP) address is shared by all the devices in the cluster. In an active/passive configuration, you configure the VIP address.
- You can deploy active/passive clustering only within the same IP subnet.

Configuring an Active/Passive Cluster

You use the primary node admin GUI to create the cluster and add members. The primary node is added as part of the cluster creation operation. When you add members, you are prompted for settings unique to the member, such as the name and IP address configuration for the internal and external interfaces. A few additional settings are also unique, namely the management port and VLAN port settings, so you add these manually after the add node procedure that follows, but before the join cluster operation.

Note: If IPv6 is required, then configure both the nodes with IPv6 settings before creating the cluster.

To create a cluster and add members:

1. Select **System > Clustering > Create Cluster** and enter a name for the cluster, a cluster password, and a name for this node, such as Node-X.

   You need to enter the password again when specifying additional nodes to join the cluster. All nodes in the cluster use this password to communicate.

Figure 13 shows the Create New Cluster page.

Figure 13 Create New Cluster Page

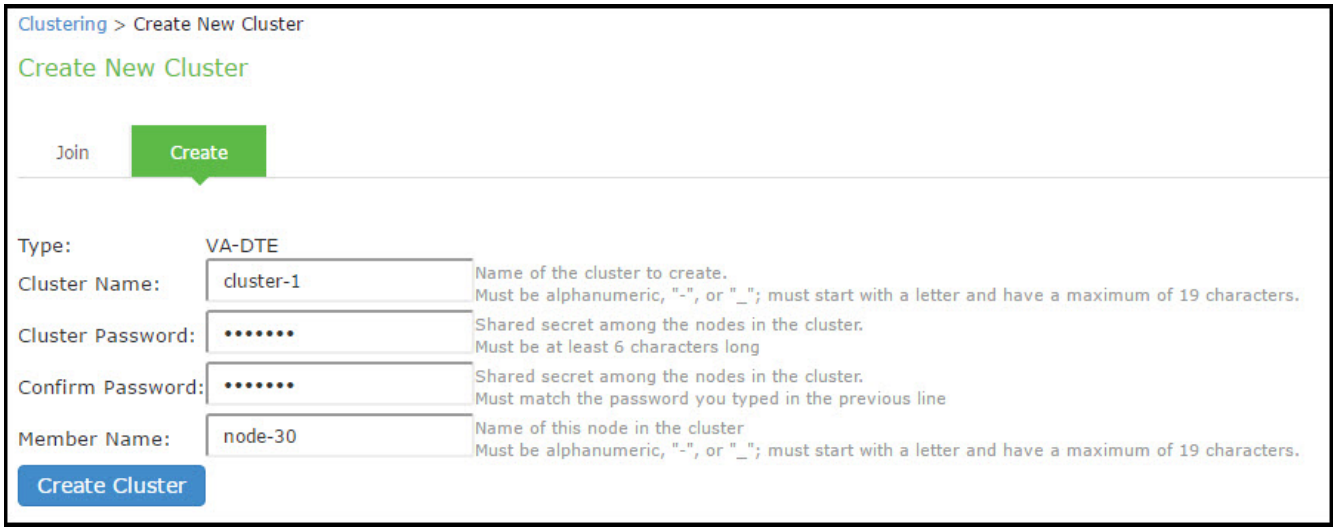

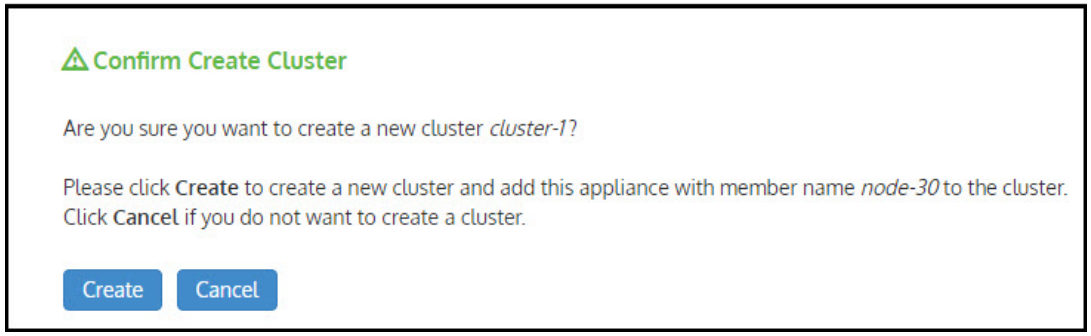

2. Click **Create Cluster.** When prompted to confirm the cluster creation, click **Create**. After the device initializes the cluster, the Clustering page displays the Status and Properties tabs.

3. Click **Properties** and select **Active/Passive configuration.**

Figure 14 shows the Clustering page for Pulse Connect Secure.

Figure 14 Clustering Page- Active/Passive Configuration

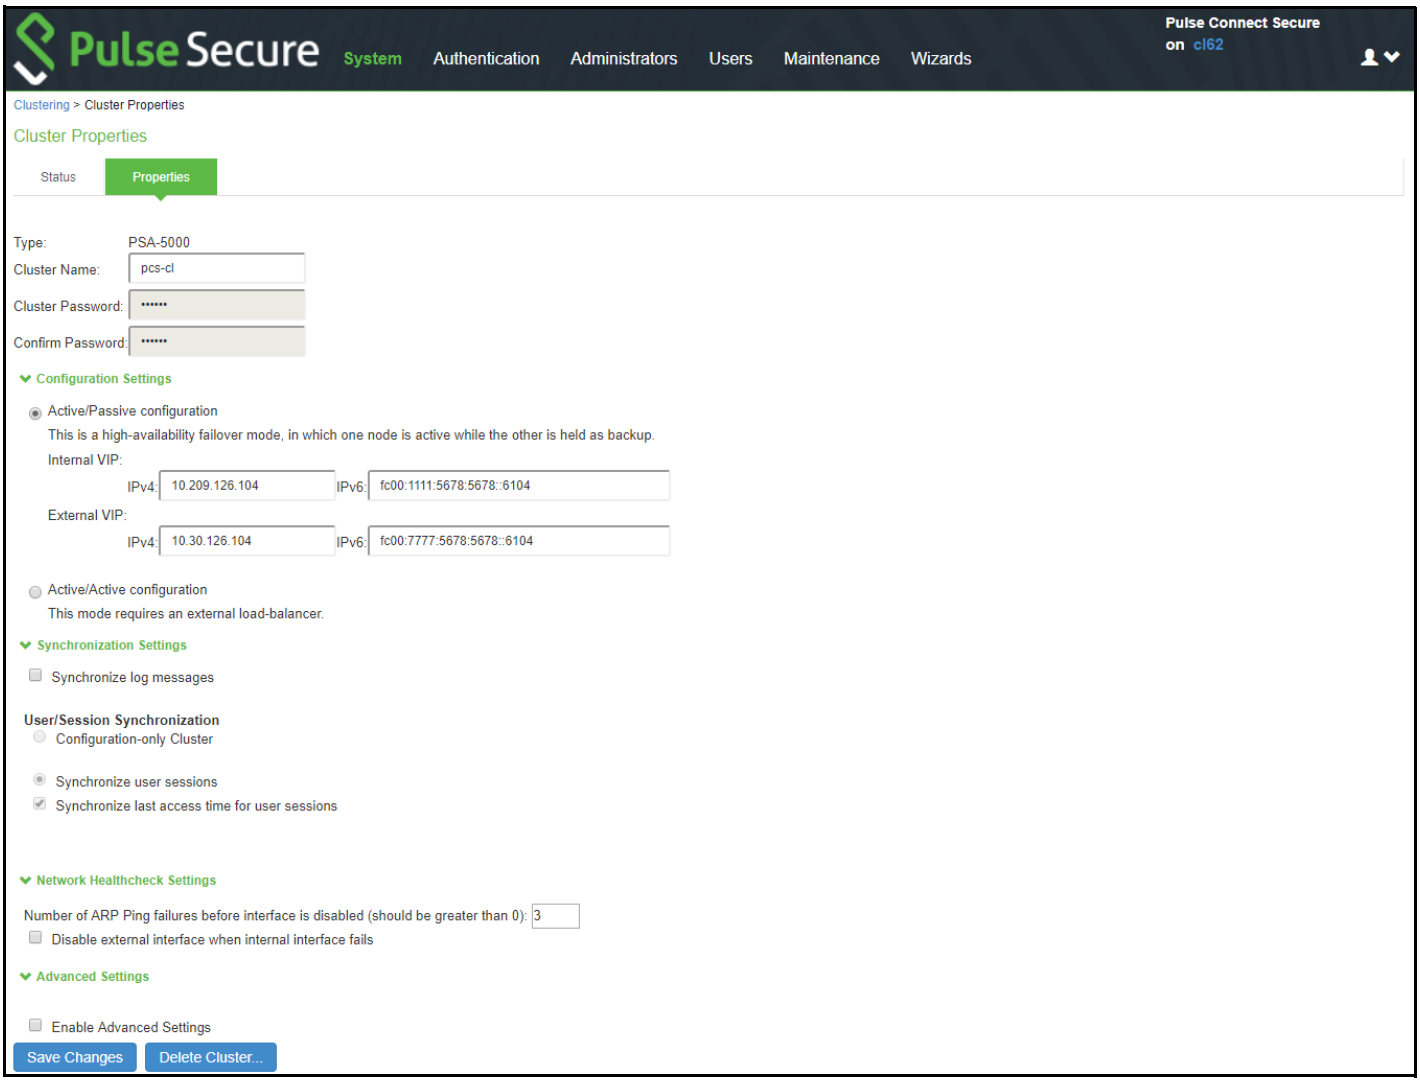

4. Complete the configuration as described in Table 3.

Table 3 Clustering Property Settings

Settings	Guidelines
Cluster Name	Specifies a name to identify the cluster.
Configuration Settings	
Active/Passive configuration	Select this option to run a cluster pair in active/passive mode. Then, specify an internal VIP (virtual IP address) and an external VIP if the external port is enabled.

Settings	Guidelines
Active/Active configuration	Select this option to run a cluster pair in active/active mode. Active/Active runs a cluster of two or more nodes in active/active mode using an external load balancer. To change a two-unit active/passive cluster to an active/active cluster with more than two nodes, first change the configuration of the two-unit cluster to active/active and then add the additional nodes.
Synchronization Settings	
Synchronize log messages	Select this option to propagate all log messages among the devices in the cluster.
User/Session Synchronization	
Configuration only cluster	Select this option to disable synchronization of session data and to replicate only configuration data and user records (for example, web bookmarks, NFS and windows shared files, terminal servers, telnet sessions, SAM, preferences, and passwords). **Note:** Enabling this option limits data transfers between the cluster nodes. User and Session specific limits are only enforced on the node and not across the cluster.
Synchronize user sessions	Select this option to synchronize all user session information (for example, instances of access to intranet services) among all the devices in the cluster.
Synchronize last access time for user sessions	Select this option to propagate the latest user access information across the cluster.

- If you configure your cluster as active/passive, the Synchronize user sessions and Synchronize last access time for user sessions options are automatically selected.
- If you select both Synchronize log messages and Synchronize user sessions check boxes, everything is replicated on the cluster nodes, including networking information. Even though networking information, including syslog and SNMP settings, can be configured per node or per cluster, all of the networking information is synchronized between nodes when these two options are set.
- If your cluster node configurations diverge because of changes made to one node while another is disabled or unavailable, the system manages the remerging of the configurations automatically, for up to 16 updates. Beyond the maximum number of allowable updates, you might need to intervene and remerge the configurations manually. In some instances, the system might be unable to remerge the configurations if there is not enough overlapping configuration information between two nodes to manage the internode communication.

For example, for a two-node cluster in which the two nodes are partitioned from each other because of a network outage, if the internal network IP address of one of the nodes changes in one of the partitions, the two partitions are unable to rejoin, even when the network is repaired. In such a case, you must remerge the configurations manually.

Network Healthcheck Settings	
Number of ARP Ping Failures	Specify the number of ARP ping failures allowed before the internal interface is disabled.
Disable external interface when internal interface fails	Select this option to disable the external interface of the device if the internal interface fails.
Advanced Settings	
Enable Advanced Settings	Select the Advanced Settings check box to specify the timeouts for the underlying cluster system. Do not change any values under this setting unless instructed to do so by Pulse Secure Technical Support.

Settings	Guidelines
Network Type	Select the appropriate network type. Network type selection controls the timeouts used by the underlying cluster system. Change this value only when you observe repeated cluster partitions that may be related to long network delays or significant load in any of the cluster nodes. A non-default network type cannot be used in conjunction with non-default timeout multipliers. If a non-default network type is picked, the timeout multiplier will silently get reset to the default value.
Timeout Multiplier	Default cluster timeouts have been picked to be optimal for typical cluster installations. Administrators have the ability to adjust the cluster timeouts over a linear scale of 1-20. Smaller timeouts result in faster failure detection. Larger timeouts minimize the risk of cluster splits during transient network glitches. The system can be instructed to pick a reasonable default for the current cluster configuration by specifying a value of 0. A non-default timeout multiplier can only be used in conjunction with the default network type. If a non-default network type is picked, the timeout multiplier will silently get reset to the default value.

5. Click **Save Changes.** After Connect Secure initializes the active/passive cluster, the Clustering page displays the **Status** and **Properties** tabs.
6. Click **Add Members** to specify additional cluster nodes.

Figure 15 shows the page for Pulse Connect Secure.

Figure 15 Add Cluster Member Page

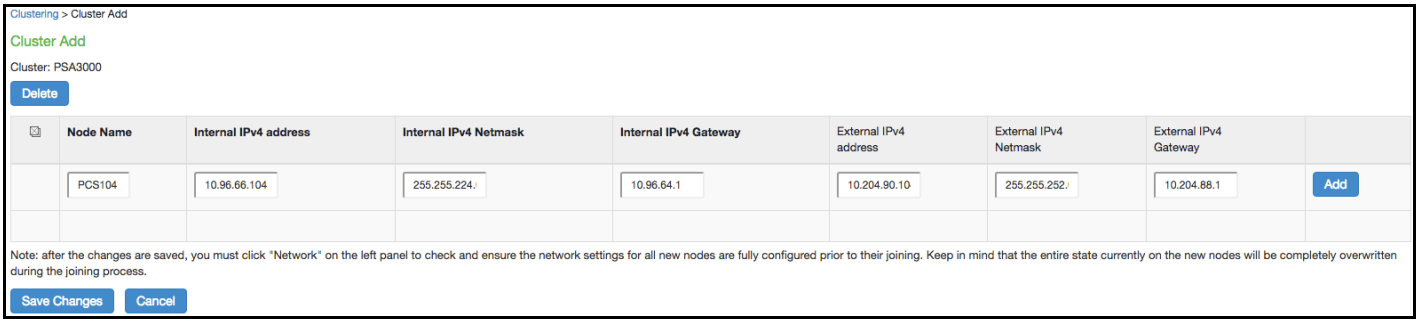

7. Click **Save Changes**.
8. Select **System > Network > Management Port > Settings** and configure the management port IPv4 and IPv6 (if configured) of node-2.

Figure 16 Configuring Management Port

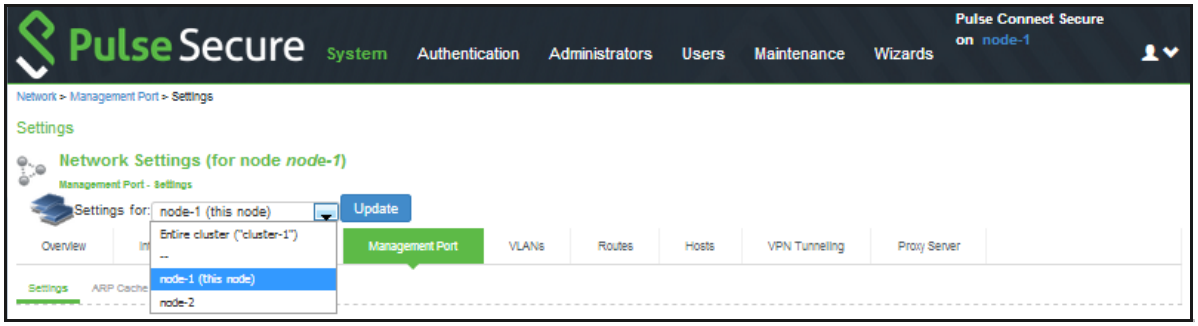

9. If a license server needs to be configured on both the nodes of a cluster, then perform the following steps:

   a. Navigate to **Configuration >Licensing > Configure Server.**

   b. Select the setting for **Entire cluster.**

   c. Configure the **License server IP** and preferred network.

   d. Click **Save Changes**.

Figure 17 Configuring License Server for Entire Cluster

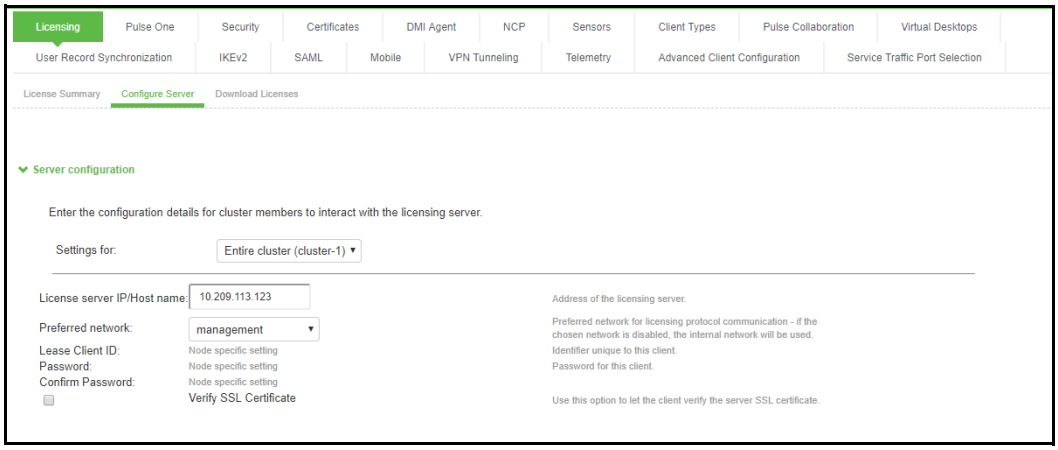

   e. Now, select the settings for node-wise and provide **Lease Client ID, Password and Çonfirm Password** for each node.

Figure 18 Node-wise Server Configuration

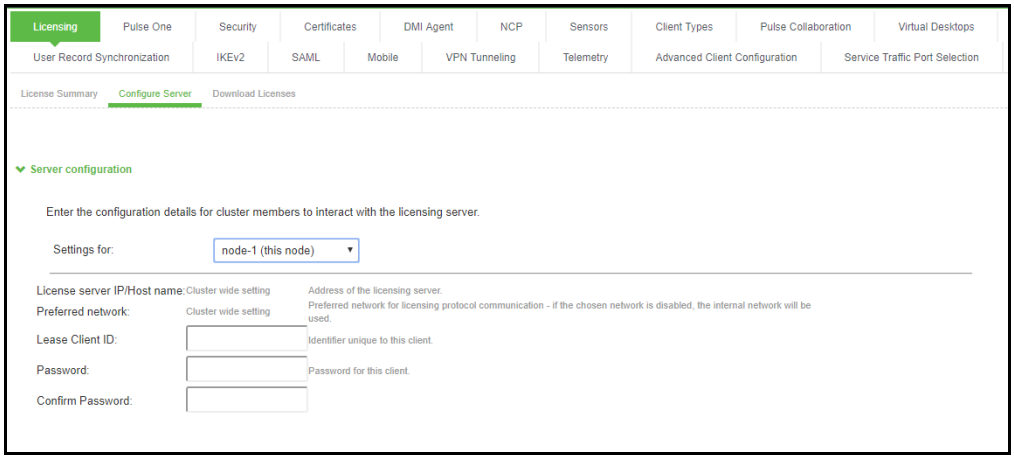

Joining Nodes to the Cluster

The primary node joins the cluster as part of the creation process.

To join additional nodes to the cluster:

1. From an existing cluster member, select the **System > Clustering > Cluster Statu**s tab and specify the node you want to add to the cluster.
2. From the admin console of the node you want to add to a cluster:
   a. Select the **System > Clustering > Join** tab and enter:
      - The name of the cluster to join
      - The cluster password you specified when defining the cluster
      - The IP address of an active cluster member
   b. Click **Join Cluster**. When prompted to confirm joining the cluster, click Join.

While the new node synchronizes its state with the existing cluster member, each node's status indicates Enabled, Enabled, Transitioning, or Enabled, Unreachable.

When the node finishes joining the cluster, its Clustering page shows the Status and Properties tabs. After the node joins the cluster, you might need to sign in again.

Verifying

Purpose	Verifying the configuration on **System > Clustering > Cluster Status** page.
Action	Select **System > Clustering > Cluster Status.**

Figure 19 shows the status on the Clustering page for Pulse Connect Secure.

Figure 19 Clustering Page -Status

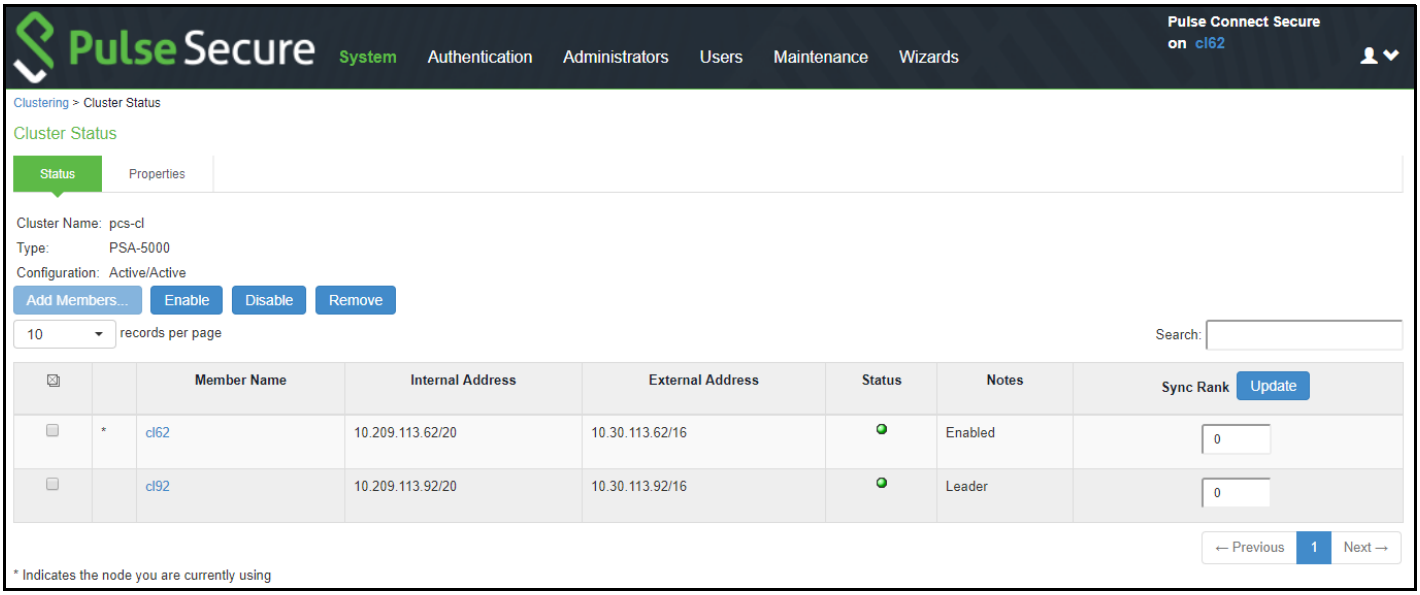

Table 4 describes the information displayed on the Status tab and the various management tasks you can perform, including disabling, enabling, and removing a node from a cluster.

Table 4 Clustering Status

GUI Element	Description
Status Information labels	Displays the cluster name, type, configuration, internal VIP, and external VIP for an active/passive cluster.
Add Members button	Click this button to specify a node you intend to add to the cluster. You can add multiple nodes at the same time.
Enable button	Click this button to add a node that was previously disabled. When you add a node, all state information is synchronized on the node.
Disable button	Click this button to disable a node within the cluster. The node retains awareness of the cluster but does not participate in state synchronizations or receive user requests unless members sign in to the node, directly.
Remove button	Click this button to remove the selected node or nodes from the cluster. After removal, the node runs in standalone mode.
Fail-Over VIP	Click this button to failover the VIP to the other node in the active/passive cluster. Only available if cluster is configured as active/passive.
Member Name column	Lists all nodes belonging to the cluster. You can click on a node's name to modify its name and network settings.
Internal Address column	Shows the internal IP address of the cluster member using Classless Interdomain Routing (CIDR) notation.

GUI Element	Description
External Address column	Shows the external IP address of the cluster member using CIDR notation. Note that this column shows only the external IP address of the cluster leader unless you specify a different address for the node on its individual network settings page, which is accessible by clicking its name in the Member Name column. If you change the external IP address on the Network > Network Settings page, the change affects all cluster nodes.
Status column	Shows the current state of the node: **Green light, Leader** - The node is the active member of an active/active cluster and is handling user requests. **Green light/enabled** - The node is handling user requests and participating in cluster synchronization. **Yellow light/transitioning** - The node is joining the cluster. **Red light/disabled** - The node is not handling user requests or participating in cluster synchronization. **Red light/enabled, unreachable** - The node is enabled but because of a network issue, it cannot be reached. **Note:** A node's state is considered standalone when it is deployed outside of a cluster or after being removed from a cluster.
Notes column	Shows the status of the node's connection to the cluster: • **OK** - The node is actively participating in the cluster. • **Transitioning** - The node is switching from the standalone state to the enabled state. • **Unreachable** - The node is not aware of the cluster. A cluster member might be unreachable even when it's online and can be pinged. Possible reasons include: its password is incorrect, it doesn't have information about all cluster nodes, it's configured with a different group communication mode, it is running a different service package version, or the machine is turned off.
Sync Rank column	Specifies the synchronization order for nodes when a node rejoins a cluster. Accepts sync ranks from 0 (lowest rank) to 255 (highest rank). The highest rank takes precedence. If two nodes have identical sync ranks, the alphanumeric rank of the member name is used to determine precedence.
Update button	Updates the sync rank after you change the precedence of the nodes in the Sync Rank column

.Using a Load Balancer

Overview

In active/active mode, you have the option of using an external load balancer with a cluster. If you do use a load balancer, all the nodes actively handle user requests sent by the load balancer or round-robin DNS. The load balancer hosts the cluster VIP and routes user requests to a node defined in its cluster group based on source-IP routing. If a node goes off line, the load balancer adjusts the load on the active nodes. Users do not need to sign in again, however some session information entered a few seconds before the active machine went off-line, such as cookies and passwords, may not have been synchronized on the current device, in which case users may need to sign in to back-end Web servers again.

The cluster itself does not perform any automatic fail-over or load-balancing operations, but it does synchronize state data (system, user, and log data) among cluster members. When an off-line device comes back online, the load balancer adjusts the load again to distribute it among all active members. This mode provides increased throughput and performance during peak load but does not increase scalability beyond the total number of licensed users.

The system synchronizes state data on all nodes if you add or delete the host entry on the Network Settings pages. If you add or delete the host entry using the Clustering tab for a cluster member, the state data affects only the node and the system does not synchronize the data across the entire cluster.

The system hosts an HTML page that provides service status for each node in a cluster. External load balancers can check this resource to determine how to effectively distribute the load among all the cluster nodes.

Figure 20 illustrates an active/active cluster configuration in which the devices have enabled external ports.

This active/active cluster configuration is deployed behind an external load balancer. You can deploy a cluster pair or multi-unit cluster in active/active mode. User requests are directed to the cluster VIP defined on the load balancer, which routes them to the appropriate machine.

Figure 20 Active/Active Configuration

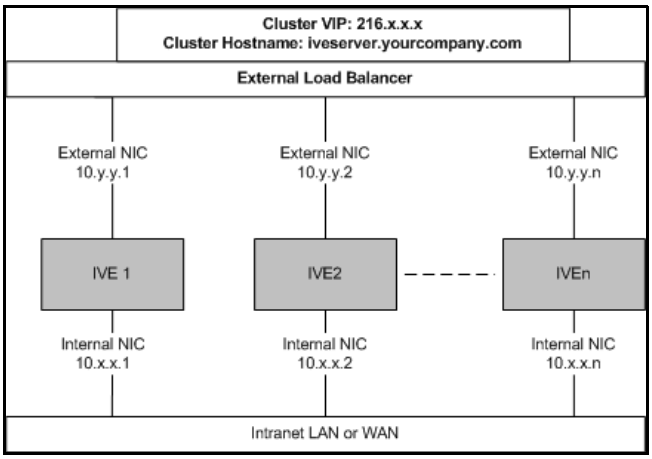

Requirements and Limitations

When choosing and configuring a load balancer for your cluster, we recommend that you ensure the load balancer:

- Supports IPsec
- Listens for traffic on multiple ports

- Can be configured to manage traffic using assigned source and destination IP addresses (not destination port)

Configuring a Load Balancer

The load balancer is configured externally.

Health Checking a Server from a Load Balancer

Purpose	The system hosts an HTML page that provides service status for each node in a cluster. External load balancers can check this resource to determine how to effectively distribute the load among all the cluster nodes.
Action	To perform the Layer 7 health check for a node:

- In a browser-Enter the URL: https://Pulse Connect Secure Controller-Hostname/dana-na//healthcheck.cgi?status=SBR

  This returns the Steel Belted Radius (SBR) status (SBR_AVAILABLE), either HTTP Status 200 OK or 500 Internal Error. If SBR_AVAILABLE is 0, the SBR is down. If SBR_AVAILBLE is 1, then SBR is up and performing transactions.

- https://Pulse Connect Secure Controller-Hostname/dana-na/healthcheck/healthcheck.cgi?status=all

  This returns either HTTP Status 200 OK or 500 Internal Error. If this returns HTTP Status 200 OK, the following additional parameters are shown:

Parameter Name	Value	Description
CPU-UTILIZATION	0-100	Specifies the CPU utilization percentage (0-100).
SWAP-UTILIZATION	integer	Specifies the swap utilization percentage of the device (0-100).
DISK-UTILIZATION	integer	Specifies the used disk space percentage (0-100).
SSL-CONNECTION-COUNT	integer	Specifies the total number of SSL connections.
USER-COUNT	integer	Specifies the total number of licensed users logged in to the device. This does not include any MAC address users or Radius users.
MAX-LICENSED-USERS-REACHED	boolean	Specifies the maximum number of licensed users reached.
VPN-TUNNEL-COUNT	integer	Specifies the number of concurrent Pulse IPSec, Network Connect and IKEv2 tunnels.
PLATFORM-LIMIT	integer	Specifies the maximum user limit on PSA hardware.
MAXIMUM-LICENSE-COUNT	integer	Specifies the maximum licenses installed directly on the PSA hardware or licenses fetched from the license server.
CLUSTER-NAME	String	Specifies the name given to the cluster. The name must be unique across the network.

The following example performs the Layer 7 health check from an external load balancer:

- GET /dana-na/healthcheck/healthcheck.cgi?status=all HTTP/1.1\r\nUser-Agent: Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 6.1; WOW64; Trident/4.0; SLCC2; .NET CLR 2.0.50727; .NET CLR 3.5.30729; .NET CLR 3.0.30729; Media Center PC 6.0; .NET4.0C; MS-RTC LM 8; .NET4.0E)\r\nHost: localhost\r\n\r\n\r\n

The concept of receive string is used by health check. The receive string is configured on the load balancer to decide whether or not to mark a node as active or inactive. It is a regular expression that checks for a value present in the response. For example, Connect Secure sends a page to the load balancer that has USER-COUNT=25 indicating that the number of active licensed users on that device is 25.

A receive string of USER-COUNT\=([0-9]|[0-9][1-9]|100); means check if USER-COUNT is between 0 and 100. In this example, 25 is between 0 and 100 and the load balancer marks the device as active and considers it for load balancing. Suppose more users log in to the device and it now sends USER-COUNT=150 to the load balancer. This value is now out of the range and the load balancer marks that device as inactive and stop sending traffic to it. Active sessions will continue to pass through the device however.

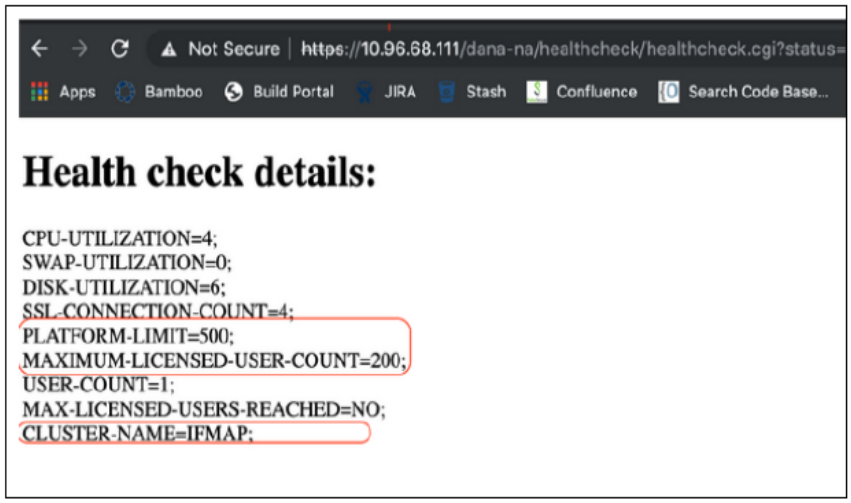

Admin Console Procedures

Creating a Cluster

To create a cluster and add members:

1. Select **System > Clustering > Create Cluster** and complete the configuration as described in Table 5.

Figure 21 shows the Create New Cluster page.

2. Click **Create Cluster**. When prompted to confirm the cluster creation, click **Create**. After the device initializes the cluster, the Clustering page displays the **Status** and **Properties** tabs.

Figure 21 Create New Cluster Page

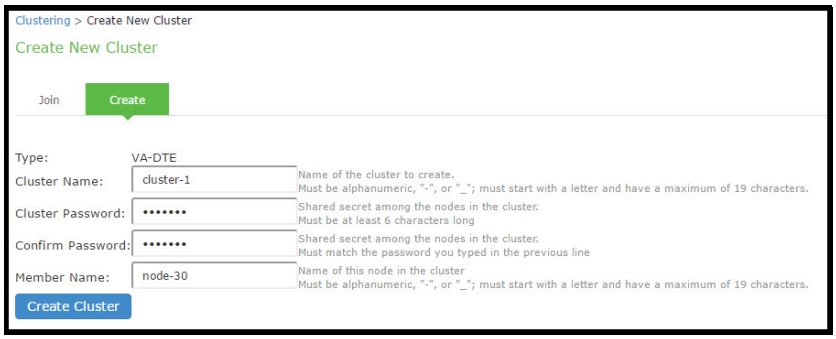

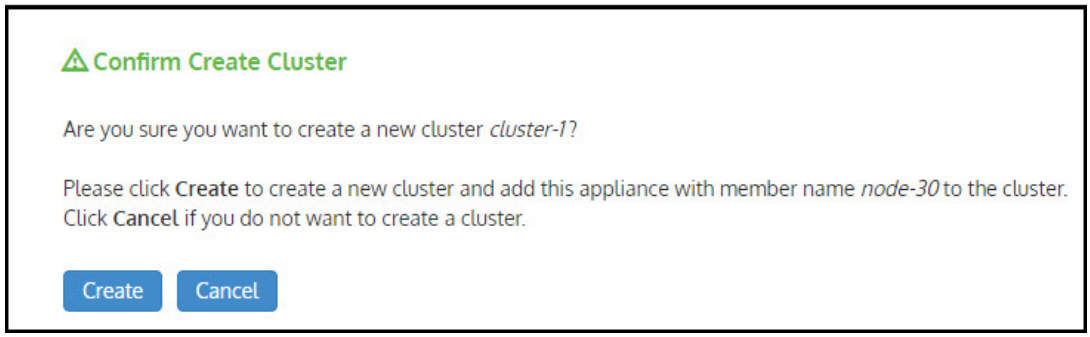

Table 5 Cluster Settings

Settings	Actions
Cluster Name	Specifies a name to identify the cluster.
Cluster Password	Specifies the cluster password. You need to enter the password again when specifying additional nodes to join the cluster. All nodes in the cluster use this password to communicate.
Confirm Password	Specifies the password that is confirmed.
Member Name	Specifies the name of the member.

Adding a Node to a Cluster Through the Admin Console

Before you can add a node to a cluster (through either the Web or the serial console), you need to make its identity known to the cluster. Note that if a node has a cluster license key, it has only a Clustering > Join tab.

To add a node to a cluster through its admin console:

1. From an existing cluster member, select **System > Clustering > Cluster Status,** and specify the node you want to add to the cluster.
2. From the admin console of the node you want to add to a cluster, select **System > Clustering > Join**, and enter:
   - The name of the cluster to join
   - The cluster password you specified when defining the cluster
   - The IP address of an active cluster member
3. Click **Join Cluster.** When you are prompted to confirm joining the cluster, click Join. After the node joins the cluster, you may need to sign in again.

Figure 22 shows the Join Cluster page.

Figure 22 Join Cluster Page

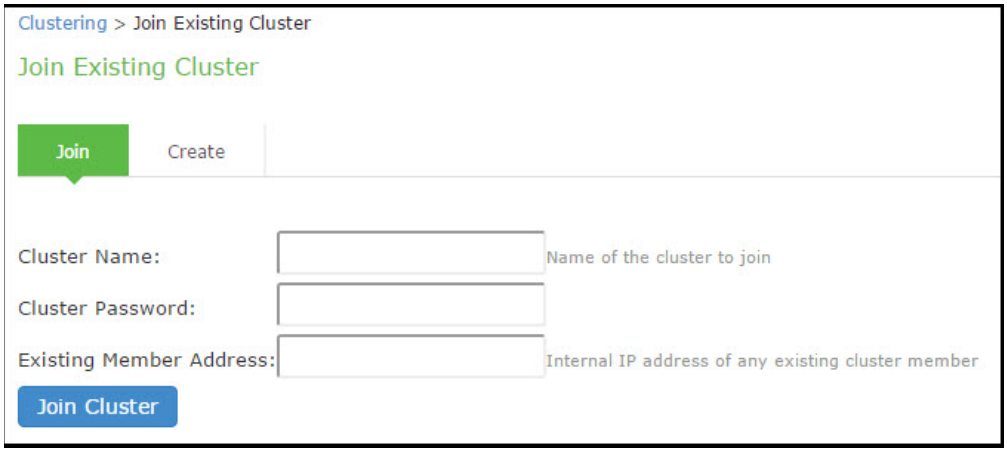

While the new node synchronizes its state with the existing cluster member, each node's status on the Status page indicates Enabled, Enabled; Transitioning; or Enabled, Unreachable.

Deleting a Cluster

If you delete a cluster, all of the nodes begin running as standalone systems.

To delete a cluster:

1. From the admin console of an active cluster member, select the **System > Clustering > Properties** page.
2. Click the **Delete Cluster** button.
3. Click **Save Changes**.

Figure 23 shows the properties for the Clustering page.

Figure 23 Clustering Page -Properties

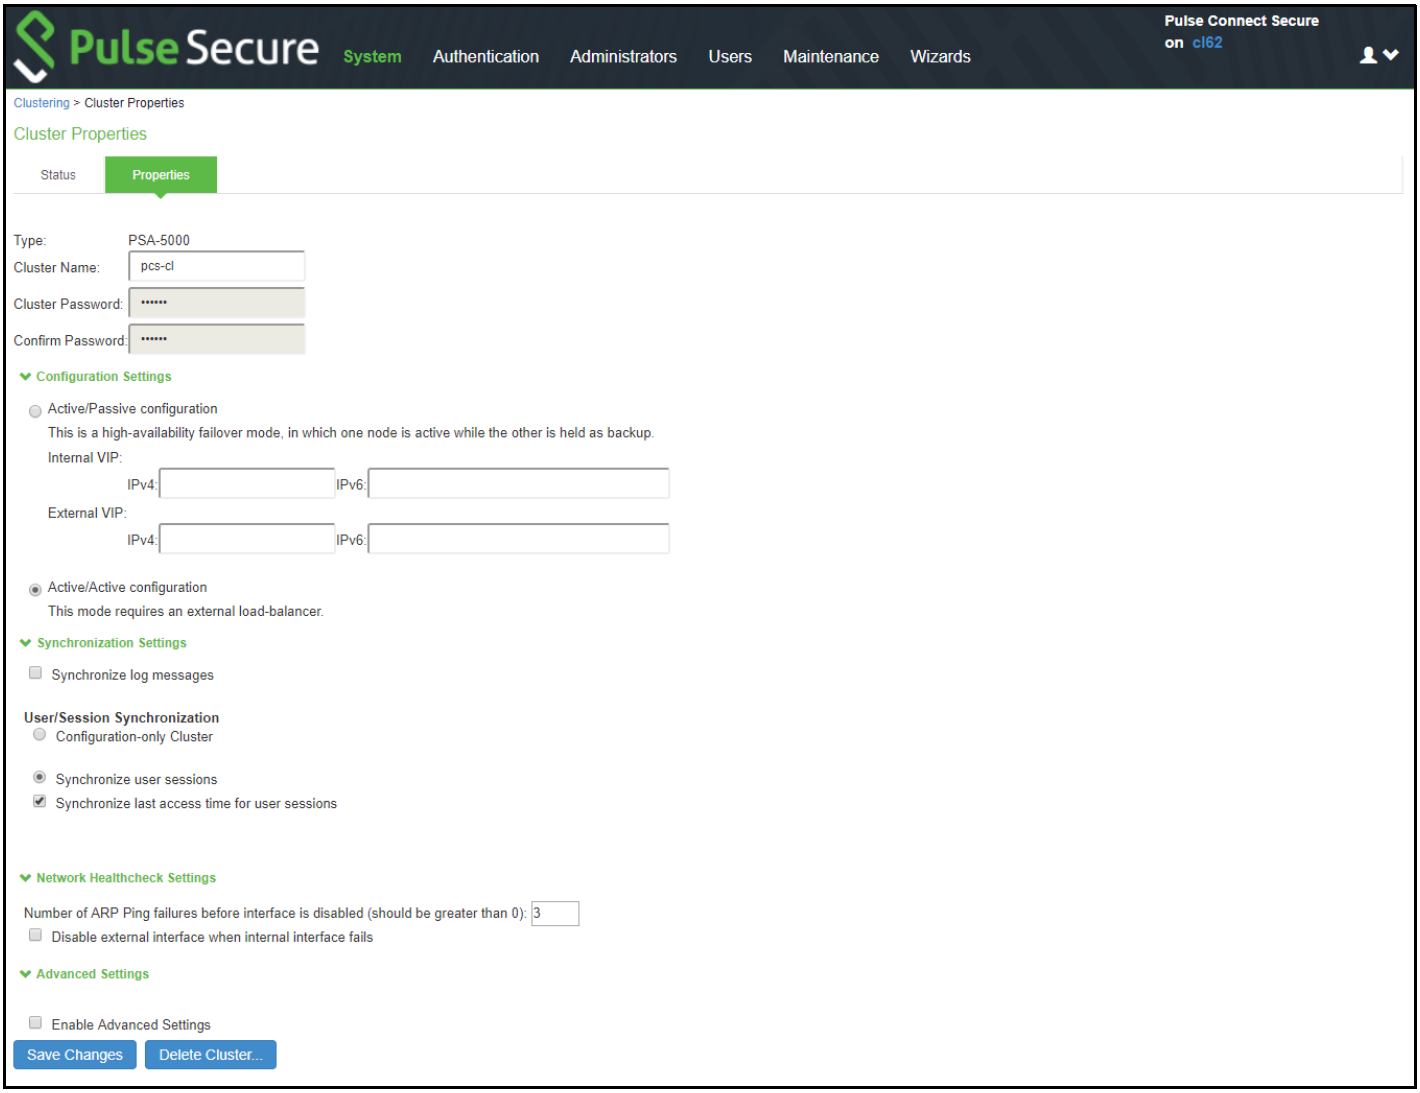

Failing Over the VIP to Another Node

In an active/passive cluster, you might need to fail the VIP to the other node, regardless of which node you are currently using.

To fail-over the VIP:

1. Select **System > Clustering > Cluster Status** from the admin console.
2. Click the **Fail-Over VIP** button to move to the other node. The Fail-Over VIP button is a toggle button, so you can move from one node to the other, regardless of which is the leader. The fail-over occurs immediately.

   Note: VIP failover does not occur when the management port fails.

Figure 24 shows the fail-over VIP option on the Clustering page.

Figure 24 Clustering Page -Status

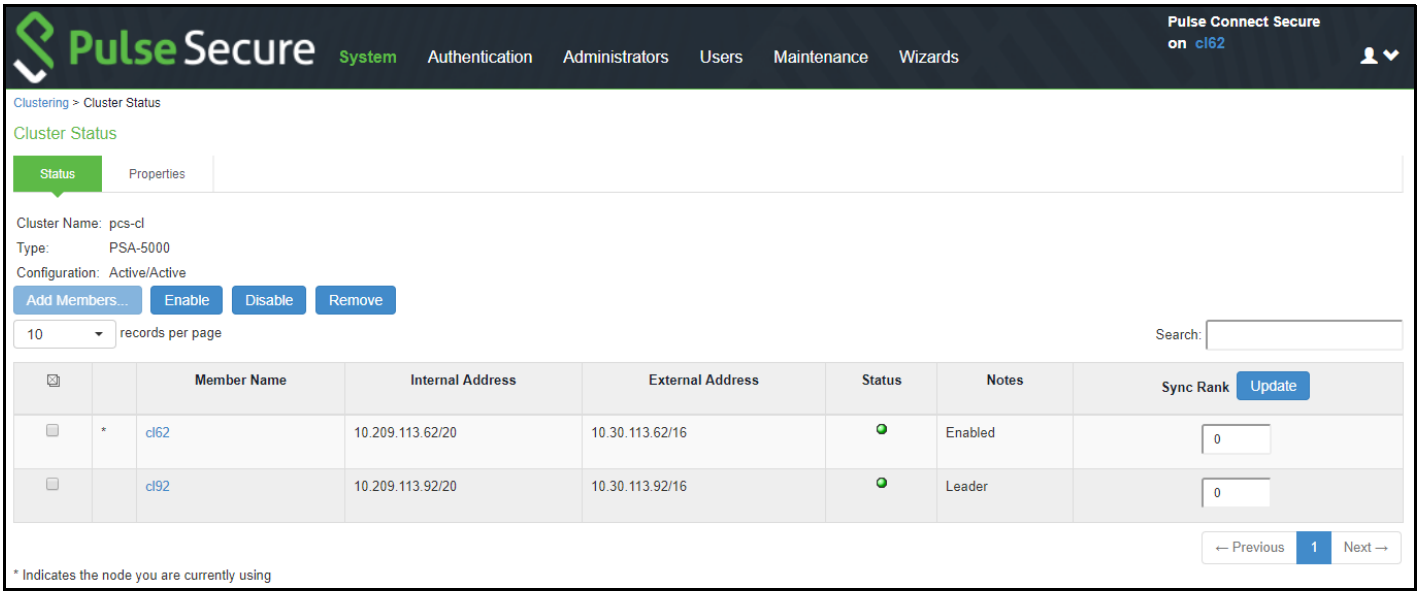

Changing the IP Address of a Cluster Node

Changing the IP address of a cluster while it belongs to a cluster is not supported. In order to change the IP address, you must first remove it from the cluster, update the IP address and then add it back.

Note: If you attempt to change the IP address of a node while it belongs to a cluster, unpredictable results might occur.

For example:

1. Select **System > Clustering > Cluster** status.
2. Select the check box for the name of the node whose IP address you want to change.
3. Click **Remove**.
4. After the node is removed, sign in to that node, change its IP address and click **Save Changes.**
5. In the main node, add the changed node to the cluster configurations.
6. Log in to the changed node and rejoin the cluster.

The following procedure is a model for changing both node IP addresses in an active/passive cluster:

1. Select **System > Clustering > Cluster** status.
2. Click **Delete Cluster.**
3. Change the IP address of each node.
4. Log in to the main node and re-create the cluster, changing it from active/active to active/passive and defining the internal and/or external VIP addresses.
5. Add the other node to the cluster configurations.

6. Log in to the passive node and add it to the cluster.

Adding Multiple Cluster Nodes

To add multiple nodes to a cluster:

Select **System > Clustering > Cluster** Status**.**

1. Click **Add Members.**
2. Enter the node name and internal IP address.
3. Modify or add the default internal netmask and internal gateway addresses, if necessary.
4. Click **Add**.

Figure 25 shows the Add Cluster Member page.

Figure 25 Add Cluster Member Page

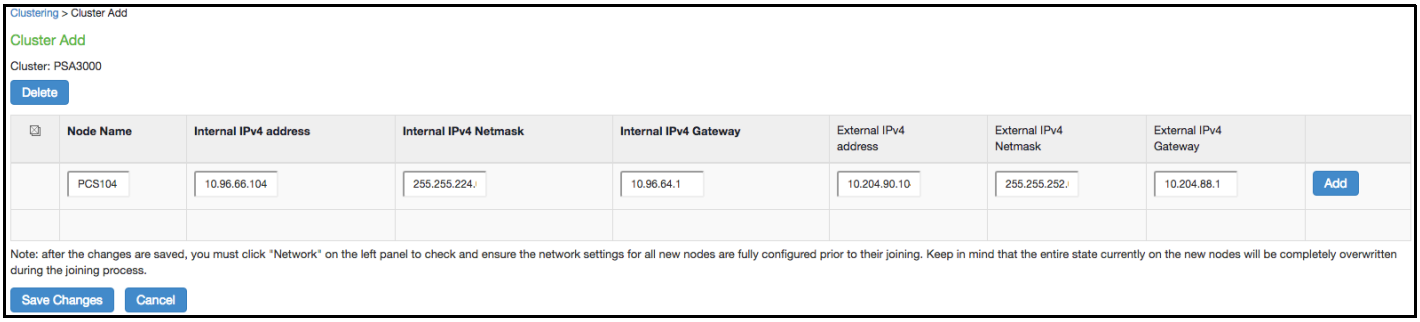

5. Repeat the process until you have added all of the nodes.
6. Click **Save Changes** to save the node configurations.

The system automatically enables the added nodes, even if they are unreachable.

Re-Adding a Node to a Cluster

With some maintenance operations, you might need to remove a node from a cluster, then re-add and re-join it to the cluster.

When a node joins a cluster, all of its node-specific settings (including network interface addresses, route tables, virtual ports, ARP caches, VLAN interface, SNMP settings) are overwritten by the corresponding configuration setting it receives from the cluster.

To populate the newly joined node with the correct node-specific settings:

1. Add the node to the cluster.
2. On any of the existing nodes in the cluster, manually configure the appropriate node-specific settings for the newly added node by selecting the node from the menu in the settings page.
3. Add the node to the cluster.

When the node joins the cluster, it receives its newly configured node-specific settings from the cluster.

Note: You configure the node-specific settings for the newly added node manually because binary import options are not useful. The only recommended binary import option into a cluster is "Import everything except network settings and licenses" from the Maintenance > Import/Export > Configuration page, which restores cluster-wide configuration (sign-in, realms, roles, resource policies etc.) from a backup binary file. Because this option skips node-specific settings, you must perform step 2 manually to populate the newly joined node with the right set of node-specific settings.

Restarting or Rebooting Cluster Nodes

When you create a cluster of two or more nodes, the clustered nodes act as a logical entity. When you reboot one of the nodes using either the serial console or the admin console, all nodes in the cluster restart or reboot.

To reboot only one node:

1. Select **System > Clustering > Status** to disable the node you want to restart or reboot within the cluster.
2. Select **Maintenance > System > Platform.**
3. Reboot the node, then enable the node within the cluster again.

The system reconciles session state with the Infranet Enforcer upon restart or cluster failover. If the Infranet Enforcer is running ScreenOS 6.0r2 or later, a Policy Secure restart or failover does not interrupt network traffic of existing sessions, as long as the restart or failover occurs within two minutes.

Figure 26 shows the System Maintenance page.

Figure 26 System Maintenance

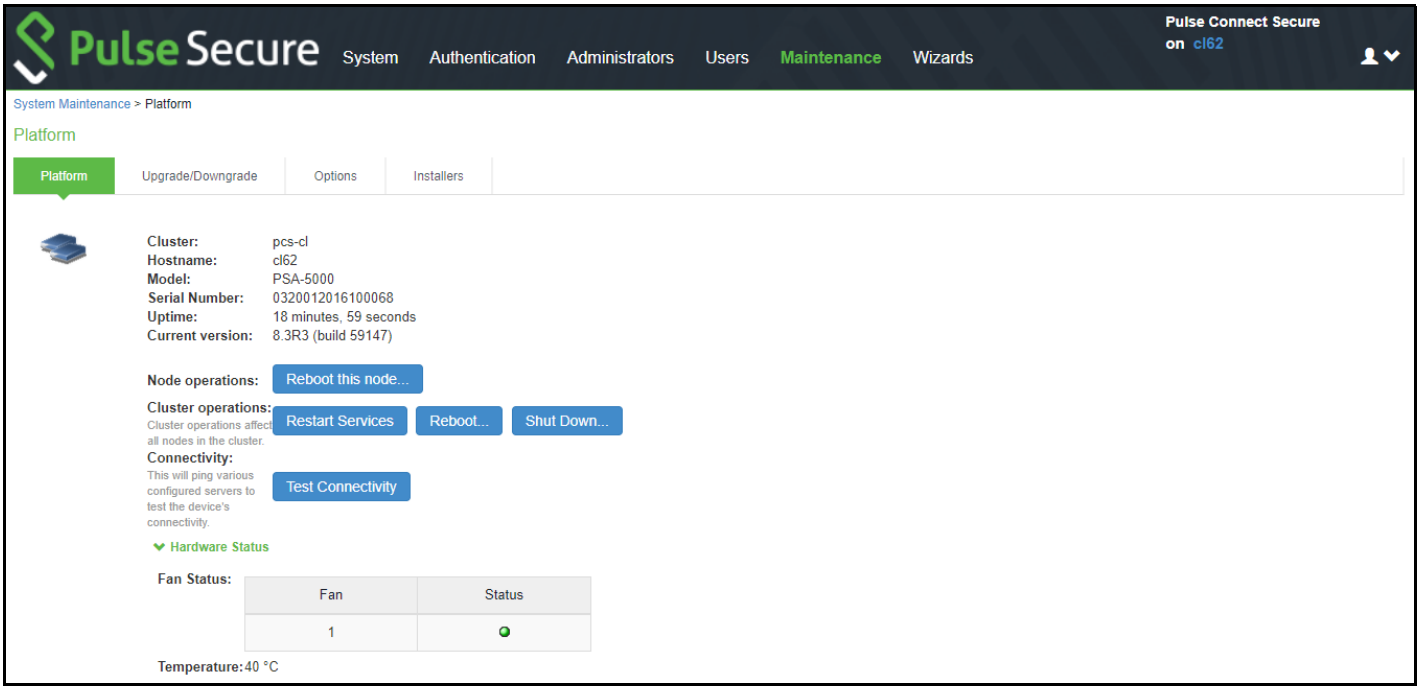

Modifying the Cluster Properties

To modify the cluster properties:

1. Select **System > Clustering > Properties.**

Figure 27 shows the properties of the Clustering page.

Figure 27 Clustering Properties Page

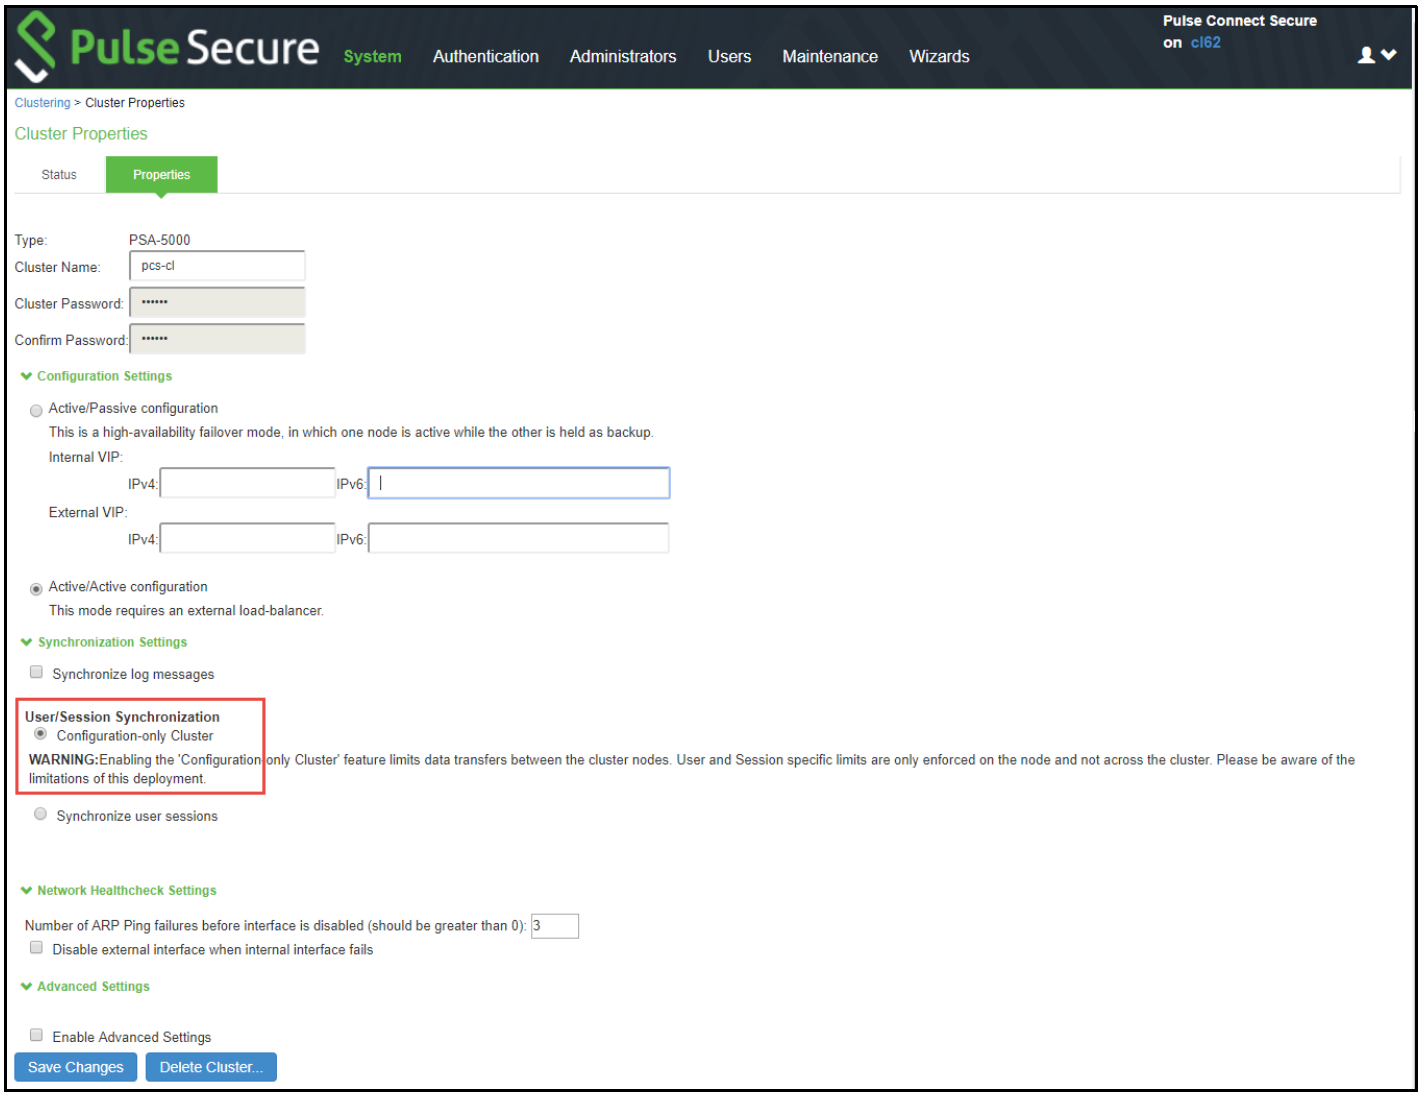

2. Complete the configuration as described in Table 6.

Table 6 Clustering Property Settings

Settings	Actions
Cluster Name	Identifies the cluster.
Configuration Settings	

Settings	Actions
Active/Passive configuration	Runs a cluster pair in active/passive mode. Then specify an internal VIP (virtual IP address) and an external VIP if the external port is enabled.
Active/Active configuration	(Default) Runs a cluster pair in active/active mode. This configuration runs a cluster of two or more nodes in active/active mode using an external load balancer. **Note:** To change a two-unit active/passive cluster to an active/active cluster with more than two nodes, first change the configuration of the two-unit cluster to active/active and then add the additional nodes.
Synchronization Settings	
Synchronize log messages	Propagates all log messages among the devices in the cluster.
User/Session Synchronization Configuration only cluster	Select this option to disable synchronization of session data and to replicate only configuration data and user records (for example, web bookmarks, NFS and windows shared files, terminal servers, telnet sessions, SAM, preferences, and passwords). **Note:** • Enabling this option limits data transfers between the cluster nodes. User and Session specific limits are only enforced on the node and not across the cluster. • Do not activate this feature when the user sessions are in progress. • Session failover is not supported in configuration only cluster mode.
Synchronize user sessions	Synchronizes all user session information (for example, instances of access to intranet services) among all the devices in the cluster.
Synchronize last access time for user sessions	Propagates the latest user access information across the cluster.

Note:

- If you select both the Synchronize log messages and Synchronize user sessions check boxes, everything is replicated on the cluster nodes, including networking information. Even though networking information, including syslog and SNMP settings, can be configured per node or per cluster, all of the networking information is synchronized between nodes when these two options are set.
- If your cluster node configurations diverge because of changes made to one node while another is disabled or unavailable, the system manages the remerging of the configurations automatically for up to 16 updates. Beyond the maximum number of allowable updates, you might need to intervene and remerge the configurations manually. In some instances, the system might be unable to remerge the configurations if there is not enough overlapping configuration information between two nodes to manage the internode communication.

For example, for a two-node cluster in which the two nodes are partitioned from each other because of a network outage, if the internal network IP address of one of the nodes changes in one of the partitions, the two partitions are unable to rejoin, even when the network is repaired. In such a case, you must remerge the configurations manually.

Network Healthcheck Settings	
Number of ARP Ping Failures	Specifies the number of ARP ping failures allowed before the internal interface is disabled.

Settings	Actions
Disable external interface when internal interface fails	Disables the external interface of the device if the internal interface fails.
Advanced Settings	
Enable Advanced Settings	Select the Advanced Settings check box to specify the timeouts for the underlying cluster system. Do not change any values under this setting unless instructed to do so by Pulse Secure Technical Support.
Network Type	Select the appropriate network type. Network type selection controls the timeouts used by the underlying cluster system. Change this value only when you observe repeated cluster partitions that may be related to long network delays or significant load in any of the cluster nodes. A non-default network type cannot be used in conjunction with non-default timeout multipliers. If a non-default network type is picked, the timeout multiplier will silently get reset to the default value.
Timeout Multiplier	Default cluster timeouts have been picked to be optimal for typical cluster installations. Administrators have the ability to adjust the cluster timeouts over a linear scale of 1-20. Smaller timeouts result in faster failure detection. Larger timeouts minimize the risk of cluster splits during transient network glitches. The system can be instructed to pick a reasonable default for the current cluster configuration by specifying a value of 0. A non-default timeout multiplier can only be used in conjunction with the default network type. If a non-default network type is picked, the timeout multiplier will silently get reset to the default value.

3. Click **Save Changes**.

Synchronizing the Cluster State

State synchronization occurs only by means of the internal network interface cards (NICs), and each cluster member is required to possess the cluster password to communicate with other members. Cluster members synchronize data when there is a state change on any member. Cluster state data is either persistent-permanently stored on the device-or transient-stored on the device only for the user's session. State data is divided into the following major categories:

- **System state** - This state is persistent and does not change often.
  - Network settings
  - Authentication server configurations
  - Authorization group configurations, such as access control list, bookmark, messaging, and application data
- **User profile** - This data can be either persistent or transient, depending on whether or not you have enabled persistent cookies and persistent password caching. If you have not enabled these features, then the data is transient and falls into the next category.
  - **User bookmarks** - persistent

  - **Persistent user cookies** - if the persistent cookies feature is enabled, the device stores user cookies for web sites that issue persistent cookies
  - **Persistent user passwords** - if the password caching feature is enabled, the user can choose to store her credentials for applications and web sites
- **User session** -This state is transient and dynamic. The user session consists of the following data:
  - The user session cookie
  - Transient user profile information, which includes cookies and passwords stored only for during the user's session
- **Monitoring state** - This persistent information consists of log messages.

  Whether you deploy a cluster in active/passive or active/active mode, the Connect Secure is responsible for synchronizing data between cluster members. The Connect Secure synchronizes all system data, user profile data, and the user session cookies immediately, so if one cluster member goes off-line, users do not need to sign in to the device again. A small amount of latency occurs when the device synchronizes user session profile and monitoring state data, so if a member goes off-line, the user may need to sign in to some back-end Web applications again and administrators may not have access to the logs on the failed machine.

  If you notice too much latency occurring on one or more nodes, you might need to change the Clustering Timeouts Settings.

  When you add the device to a cluster, the cluster leader does not send log messages to the new member. Log messages are also not synchronized between cluster members when one member restarts its services or when an offline machine comes back online. Once all machines are online, however, log messages are synchronized.

Note: If you are running an active/active cluster, you must not allow the cluster to switch to active/passive mode unless the active/active and active/passive clusters share compatible spread timeout settings.

You may also configure synchronization settings to improve performance:

- **Specify the synchronization protocol** - When running three or more devices in a multi-unit or multi-site cluster, you can choose to use the synchronization protocol (Unicast, Multicast, or Broadcast) that best suits your network topology.
- **Synchronize log messages** - Log messages may create a huge payload on the network and affect cluster performance. This option is disabled by default.
- **Synchronize user sessions** - This option synchronizes all user session information (instances of access to intranet services, for example) among all devices in the cluster.

You must select this option if your cluster is an IF-MAP client. If you do not select this option, your IF-MAP client may not work as expected.

- Synchronize last access time for user sessions-This option allows you to propagate user access information in the cluster. If this option is the sole synchronization item among the cluster nodes, you can significantly reduce CPU impact among the cluster devices.

Note:

- If you configure your cluster as active/passive, the Synchronize user sessions and Synchronize last access time for user sessions options are automatically checked.

- If you select both the both Synchronize log messages and Synchronize user sessions check boxes, everything is replicated on the cluster nodes, including networking information. Even though networking information, including syslog and SNMP settings, can be configured per node or per cluster, all of the networking information is synchronized between nodes when these two options are set.
- If your cluster node configurations have diverged due to changes made to one node while another is disabled or unavailable, the devices manage the remerging of the configurations automatically, for up to 16 updates. Beyond the maximum number of allowable updates, you may need to intervene and remerge the configurations manually. In some instances, the devices may be unable to remerge the configurations if there is not enough overlapping configuration information between two nodes to manage the internode communication.

For example, given a two-node cluster in which the two nodes are partitioned from each other because of a network outage, if the internal network IP address of one of the nodes gets changed in one of the partitions, the two partitions are unable to rejoin, even when the network is repaired. In such a case, you must manually remerge the configurations.

General Cluster Maintenance

Managing Network Settings for Cluster Nodes

To modify the network settings for a cluster or each individual node in a cluster, click System > Network. You can make your changes on the Network Settings pages. After you create a cluster, these pages provide a drop-down list from which you can select the entire cluster or a specific node to modify. When you save changes on a Network page, the settings are saved for the specified cluster or cluster node. If you change network settings for an entire cluster, they propagate to every node in the cluster.

You can access a node-specific Network page by clicking **System > Clustering > Cluster Status** on the node's name in the Member Name column.

Upgrading Clustered Nodes

The Connect Secure offers the ability to easily upgrade every node in a cluster. You simply install a newer service package on one node and, once the installation completes and the node reboots, the node pushes the service package to all nodes in the cluster.

Upgrading the Cluster Service Package

Install a newer service package on one cluster node only. When the installation process completes and the cluster node reboots, it instructs the other nodes to upgrade.

Migrating Cluster Configurations to a Replacement Cluster

To migrate system and user configurations from a Connect Secure cluster (C1) to a replacement cluster (C2) using different Connect Secure devices:

1. Export the system and user configuration from C1's primary node (PN1).

   Note the following information:

   - Cluster name

- Cluster password
- Name of the node where the export was done (PN1)
- Internal IP address of PN1
- Internal network mask of PN1
- Internal network gateway of PN1
- Name of all other nodes in the C1 cluster, including their internal network IP address, network masks and gateways

2. Shut down all Connect Secure devices in cluster C1.
3. Power on one of the new servers (must be running software release 6.1R1 or later) that is part of cluster C2 and is on the same network to which PN1 was attached. This Pulse server device is called PN2 for the remainder of these steps.
4. When prompted, configure the internal network settings of PN2 to the same internal network settings of PN1.
5. Install the new primary license on PN2.
6. From the admin GUI on PN2, select **System > Clustering> Create Cluster**. Create the cluster C2 using the same cluster name and cluster password that were in use at cluster C1. Node PN2 must also be assigned the same node name as PN1.
7. Open the cluster status page and add the remaining nodes to the cluster configuration. Nodes being added must be assigned the same names that existed in original cluster C1. The internal network settings of the newly added nodes must also match the corresponding settings in the original cluster C1.

   Note: Do not join the newly added nodes to cluster C2 yet.
8. Import the data exported from PN1 into PN2.
9. When importing the system configuration, select the option Import everything (except Device Certificate(s)).
10. Power on the remaining new Pulse Connect Secure devices assigned to cluster C2. Configure the bare minimal internal network settings needed to bring up the machine. The network settings must match what has already been configured on node PN2.

    Note: Do not do make any other configuration changes on these machines as they will be lost when these machines join the cluster. Do not add licenses on these machines yet.
11. Join the Pulse Connect Secure to cluster C2 and wait for the cluster status to stabilize.
12. Install the CL licenses on the newly joined nodes.

Configuring the External VIP for An Active/Passive Cluster

To add an external VIP to an existing A/P cluster:

1. Create an A/P cluster with only the internal port configured.
2. Select **System > Clustering > Clustering Properties** and add the internal VIP.
3. Select **System > Network > External Port.**
4. From the **Settings for** menu, select "**entire cluster".**
5. Add the **Netmask** and **Default Gateway** but leave the external port disabled.
6. For each node, select **System > Network > External Port** and configure the external port IP address but leave the external port disabled.
7. Add the external cluster VIP.
8. Select **System > Network > External Port,** select "entire cluster" from the Settings for menu and enable the external port.

Monitoring Clusters

You can monitor clusters using the standard logging tools provided by the Pulse Connect Secure. In particular, you can use several cluster-specific SNMP traps to monitor events that occur on your cluster nodes, such as:

- External interface down
- Internal interface down
- Disabled node
- Changed virtual IP (VIP)
- Deleted cluster node (cluster stop)

Note: Generally, it is desirable to configure your SNMP traps on a cluster-wide basis, so that any given cluster node can send its generated traps to the right target. Setting up cluster-wide configuration for the traps is particularly important when you also use a load balancer, because you may not know which node is responsible for a specific operation. In that case, the load balancer may independently determine which cluster node can manage an administrative session.

You can use SNMP traps that are included in the Pulse Secure Standard MIB to monitor these events. These traps include:

- **iveNetExternalInterfaceDownTrap** - Supplies type of event that brought down the external interface.
- **iveNetInternalInterfaceDownTrap** - Supplies type of event that brought down the internal interface.
- **iveClusterDisableNodeTrap** - Supplies the cluster name on which nodes have been disabled, along with a space separated list of disabled node names.
- **iveClusterChangedVIPTrap** - Supplies the type of the VIP, whether external or internal, and its value before and after the change.
- **iveClusterDelete** - Supplies the name of the cluster node on which the cluster delete event was initiated.

These traps are always enabled and available in the MIB. You cannot disable the traps.

Troubleshooting Clusters

When you have problems with cluster communication, you may be directed by your Pulse Secure Support representative to use the cluster node troubleshooting tools.

To use the cluster node troubleshooting tools:

From the admin console, select **Maintenance > Troubleshooting > Monitoring > Node Monitor, in Maintenance > Troubleshooting > Clustering Network Connectivity**, and in **Maintenance > Troubleshooting > Clustering Group Communication.**

You can use a built-in feature on the clustering Status page to identify the status of each cluster node. Pause the mouse pointer over the Status light icon and the system displays a tool tip containing a hexadecimal number. The hexadecimal number is a snapshot of the status of the Pulse Connect Secure. It is a bit mask indicating a number of states as shown in Table 7.

Table 7 Cluster Status

Value	Meaning
0x000001	Pulse Connect Secure is in standalone mode.
0x000002	Pulse Connect Secure is in cluster disabled state.
0x000004	Pulse Connect Secure is in cluster enabled state.
0x000008	Unable to communicate (because it is offline, has wrong password, has different cluster definition, different version, or a related problem).
0x00002000	The node owns the VIPs (on) or not (off).
0x000100	Pulse Connect Secure is syncing state from another Pulse Connect Secure (initial syncing phase).
0x000200	Pulse Connect Secure is transitioning from one state to another.
0x00020000	The group communication subsystems at the local and remote nodes are disconnected from each other.
0x00040000	Management interface (mgt0) appears disconnected.
0x00080000	Management gateway is unreachable for ARP ping.
0x000800	Pulse Connect Secure int0 appears disconnected (no carrier).
0x001000	This node is configured to be a cluster member.
0x002000	Pulse Connect Secure is syncing its state to another Pulse Connect Secure that is joining.
0x004000	Initial Synchronization as master or slave is taking place.
0x008000	This Pulse Connect Secure is the leader of the cluster.
0x010000	The group communication subsystem is functional.
0x020000	The gateway on int0 is unreachable for ARP pings (see log file).
0x040000	The gateway on int1 is unreachable for ARP pings (see log file).
0x080000	Leader election is taking place.

Value	Meaning
0x100000	Server life cycle process (dsmon) is busy.
0x200000	System performs post state synchronization activities.
0x30004	• "The group communication subsystem is functional. • The gateway on int0 is unreachable for ARP pings (see log file). • Pulse Connect Secure is in cluster enabled state.
0x80000000	Cluster keystore or security world has not been associated with the FIPS card.

Each code, as you see it in the Pulse Connect Secure, may relate specifically to one state. However, each code may represent a combination of states, and so the actual code does not appear in Table 7. Instead, the code you see in the Pulse Connect Secure is the sum of several of the hexadecimal numbers shown in Table 7. You will need to factor out the codes, as in the following example:

- 0x38004 - The right-most digit (4) in this hexadecimal number corresponds to:
  - 0x000004 - The Pulse Connect Secure is in cluster enabled state.
- 0x038004 - The digit in the fourth position from the right (8) corresponds to:
  - 0x008000 - This Pulse Connect Secure is the leader of the cluster.
- 0x38004 - The left-most digit (3) in this hexadecimal number does not exist in the table, which indicates that it corresponds to the sum of two other digits, in this case, 1 and 2, as shown in the following codes:
  - 0x020000 - The gateway on int0 is unreachable for ARP pings (see log file).
  - 0x010000 - The group communication subsystem is functional.

"Management IP Address Differs from the Management IP Address" Error Message

If you receive the following error when joining a standalone PSA-7000C node to a cluster even though the management port is configured and enabled:

If the Management IP address (x.x.x.x) for the local system differs from the Management IP address (not entered) configured for this system in the remote system, then perform the following steps to add the node:

1. From the admin console of the primary node, select **System > Network > Management Port.**
2. Select the node to add from the drop-down list next to the "Setting for" label.
3. Enable the management port and enter the IP address, netmask and default gateway for the joining node.
4. Click **Save Changes.**
5. From the admin console of the joining node, join the cluster again.

Fail-over Transactions

In the case of a fail-over (both in active/passive and active/active configurations), all transactions currently in progress (such as telnet or SSH sessions or large file downloads/uploads) must be restarted after the fail-over. There is no seamless fail-over for on-going transactions using sockets except for HTTP requests or non-stateful connections.

Using the Serial Console for Cluster Administration

If you are adding a factory-set device to a cluster, we recommend that you use the serial console, which enables you to join an existing cluster during the initialization process by entering minimal information. When a node joins a cluster, it receives the cluster state settings, which overwrite all settings on a device with an existing configuration and provide new machines with the required preliminary information. You can also use the serial console to disable the node. If the node is in a synchronization state, you cannot access its admin console. Therefore, if you need to upgrade or reboot the node, for example, you must first disable the node from a cluster through its serial console.

- "Joining a Node to a Cluster Using Its Serial Console" on page 44
- "Disabling a Clustered Node Using Its Serial Console" on page 45
- "Restarting or Rebooting Cluster Nodes Using Its Serial Console" on page 45

Joining a Node to a Cluster Using Its Serial Console

Before a configured or factory-set node can join a cluster, you must make its identity known to the cluster.

Note:

- To add a node currently running as a standalone device to a cluster through its admin console, it must be running the same or a more recent version service package on the same hardware platform as the other members.
- If you add a node running an earlier version service package to a cluster, the node automatically detects the mismatch, gets the newer package from the cluster, and joins the cluster.

To add a node to a cluster through its serial console:

1. In the admin console of an existing cluster member, select **System > Clustering > Cluster Status** and specify the node to add to the cluster.
2. Connect to the serial console of the device you want to add to the cluster.
3. Reboot the device and watch its serial console. After the system software starts, a message appears stating that the device is about to boot as a standalone node and to press the Tab key for clustering options. Press the Tab key as soon as you see this option.

Note: The interval to press the Tab key is five seconds. If the device begins to boot in standalone mode, wait for it to finish and then reboot again.

4. Enter the number instructing the node to join an existing cluster.
5. Enter the requested information, including:
   - The internal IP address of an active member in the cluster

- The cluster password, which is the password you entered when defining the cluster
- The name of the device to add
- The internal IP address of the device to add
- The netmask of the device to add
- The gateway of the device to add

The active cluster member verifies the cluster password and that the new device's name and IP address match what you specified in the admin console. If the credentials are valid, the active member copies all of its state data to the new cluster member, including certificate, user, and system data.

6. Enter the number instructing the node to continue the join cluster operation. When you see a message confirming that the device has joined the cluster, select System > Clustering > Cluster Status in the admin console of any active cluster member to confirm that the new member's Status is green, indicating that the node is now an enabled node of the cluster (status is green).

Disabling a Clustered Node Using Its Serial Console

To disable a node within a cluster using its serial console:

1. Connect to the serial console of the device you want to disable within the cluster.
2. Enter the number that corresponds to the System Operations option.
3. Enter the number that corresponds to the Disable Node option.
4. Enter y when the serial console prompts you to confirm that you want to disable the node.
5. Verify that the node has been disabled (status is red) within the cluster by selecting System > Clustering > Status in the admin console of any active cluster member.

Restarting or Rebooting Cluster Nodes Using Its Serial Console

When you create a cluster of two or more nodes, the clustered nodes act as a logical entity. When you reboot one of the nodes using either the serial console or the admin console, all nodes in the cluster restart or reboot.

To reboot only one node:

1. Connect to the serial console of the device you want to disable within the cluster.
2. Enter the number that corresponds to the **System Operations** option.
3. Select **System > Clustering > Status** to disable the node you want to restart or reboot within the cluster.
4. Under system operations select the appropriate menu option <Reboot this device>, <Shutdown this device>, or <Restart Services>.
5. Reboot the node, then enable the node within the cluster again.

The system reconciles session state with the Infranet Enforcer upon restart or cluster failover. If the Infranet Enforcer is running ScreenOS 6.0r2 or later, a Policy Secure restart or failover does not interrupt network traffic of existing sessions, as long as the restart or failover occurs within two minutes.

Monitoring Cluster Nodes

If you have a problem with a cluster, a Pulse Secure Support representative may ask you to create a snapshot that includes node monitoring statistics to assist with debugging the cluster problem. When you enable the node monitor on the Maintenance > Troubleshooting > Monitoring > Node Monitor tab, the Pulse Connect Secure captures certain statistics specific to the cluster nodes on your system. Using the snapshot that results, the support team can identify important data, such as network statistics and CPU usage statistics.

To enable node monitoring:

1. Enable the node monitor on the **Maintenance > Troubleshooting > Monitoring > Node Monitor tab**
2. Enter the maximum size for the node monitor log.
3. Enter the interval, in seconds, at which node statistics are to be captured.
4. Select the **Node monitoring enabled** check box to start monitoring cluster nodes.
5. For **Maximum node monitor log size**, enter the maximum size (in MB) of the log file. Valid values are 1-30.
6. Specify the interval (in seconds) that defines how often nodes are to be monitored.
7. Select the commands to use to monitor the node.

   If you select **dsstatdump**, enter its parameters as well.
8. Click **Save Changes**.
9. If you want to include the node monitoring results in the system snapshot, choose **Maintenance > Troubleshooting > System Snapshot**, and select the **Include debug log** check box.
10. Take a system snapshot to retrieve the results.

Cluster Group Communication and Node Monitoring

- "Overview" on page 47
- "Configuring Cluster Network Connectivity Monitoring" on page 50
- "Configuring Cluster Node Monitoring" on page 48

Overview

If you have a problem with a cluster, a Pulse Secure Support representative might ask you to create a snapshot that includes group communication statistics to assist with debugging the cluster problem. When you enable the group communication monitor in the Group Communication tab, the system records statistics related to all of the cluster nodes on your system. As the local node communicates with other nodes in the cluster, the system captures statistics related to intra cluster communication. The Group Communication tab is displayed only when you enable clustering on your system. On a standalone system, you do not have access to the Group Communication tab.

You can also enable the cluster networking troubleshooting server on the Network Connectivity page.

Note:

- Performing excessive node monitoring can impact system performance and stability. You should only perform extensive monitoring when directed by your Pulse Secure Support representative.
- Performing log synchronization across cluster nodes can impact your system performance and stability.

Configuring Group Communication Monitoring on a Cluster

To enable group communication monitoring:

1. Enter the maximum size for the statistics log.
2. Enter the interval, in seconds, at which events are to be logged.
3. If you want to monitor all cluster nodes from the current local node, select the **Monitor all cluster nodes from this node** check box. If you do not select this option, the group communication monitor gathers statistics only for the local node.

Note: If you select the **Monitor all cluster nodes from this node** option, the cluster nodes must be able to communicate over UDP port 6543.

4. Select the **Enable group communication monitoring** check box to start the monitoring tool.
5. Click **Save Changes.**

Figure 28 shows the Troubleshooting page for group communication.

Figure 28 Troubleshooting using Group Communication

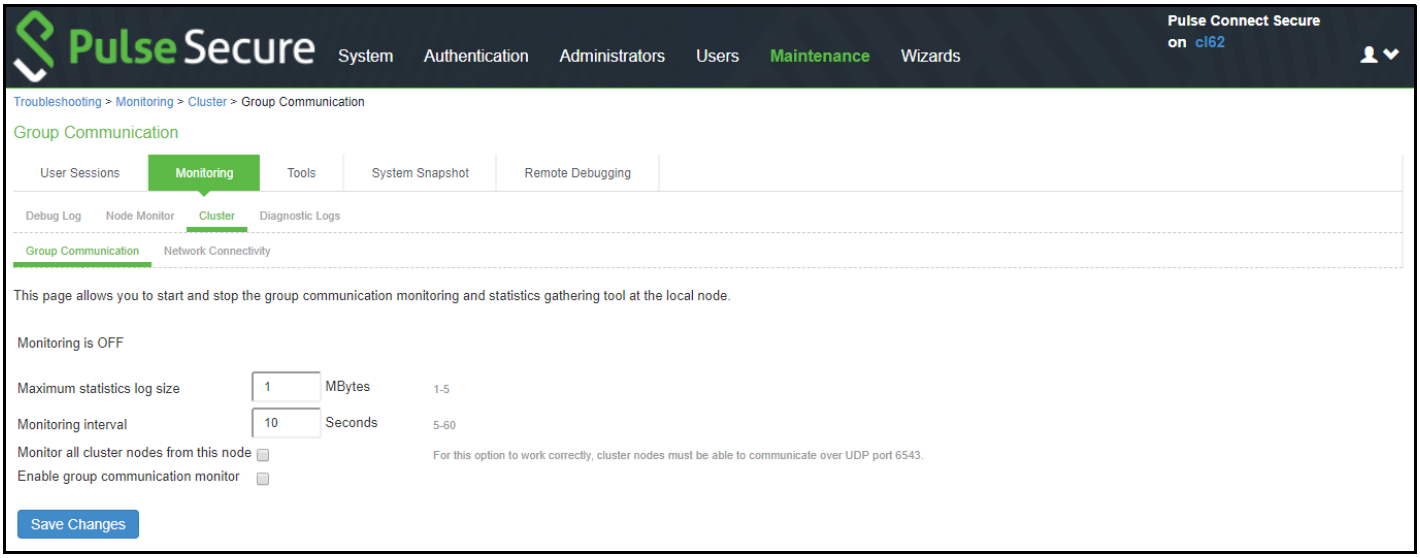

6. If you want to include the node monitoring results in the system snapshot, choose **Maintenance > Troubleshooting > System Snapshot,** and select the Include debug log check box.

7. Take a system snapshot to retrieve the results.

Configuring Cluster Node Monitoring

If you have a problem with a cluster, a Pulse Secure Support representative may ask you to create a snapshot that includes node monitoring statistics to assist with debugging the cluster problem. When you enable the node monitor on the Node Monitor tab, the IC Series device captures certain statistics specific to the cluster nodes on your system. Using the resulting snapshot, the support team can identify important data, such as network statistics and CPU usage statistics.

To enable node monitoring:

1. Select **Maintenance > Troubleshooting > Monitoring > Node Monitor** to enable the node monitor.

2. Enter the maximum size for the node monitor log.

3. Enter the interval, (in seconds) at which node statistics are to be captured.

4. Select the **Node monitoring enabled** check box to start monitoring cluster nodes.

   Figure 29 shows the Troubleshooting page for node monitoring.

   Figure 29 Troubleshooting using Node Monitor

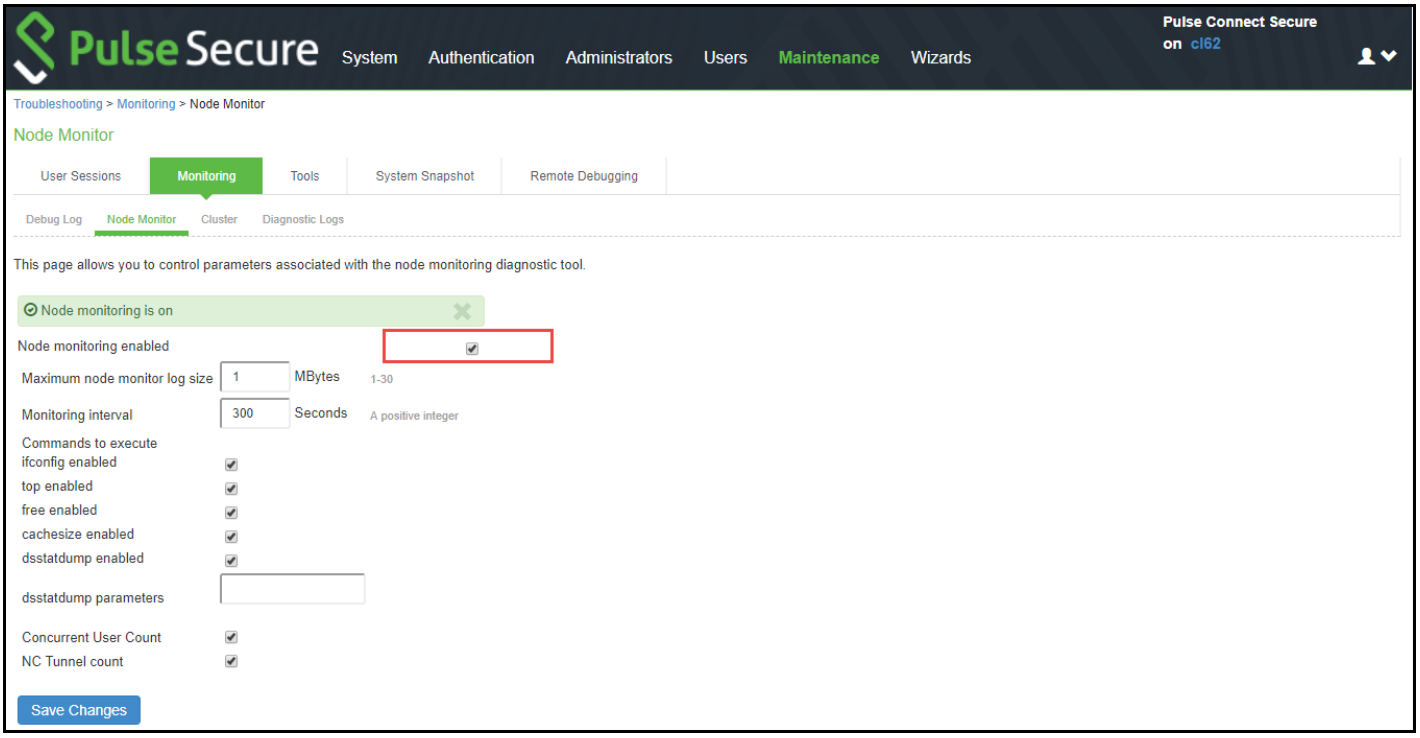

5. For **Maximum node monitor log size,** enter the maximum size (in MB) of the log file. Valid values in the range of 1 - 30.
6. Specify the interval (in seconds) that defines how often nodes are to be monitored.
7. Select the commands to use to monitor the node.

   If you select **dsstatdump,** enter its parameters as well.

   From 9.1R3 release, the "iostat" information is gathered periodically and made available as part of node monitoring in system snapshot under the "nodemon" section.
8. Click **Save Changes.**
9. To include the node monitoring results in the system snapshot, select **Maintenance > Troubleshooting > System Snapshot,** and select the Include debug log check box.
10. Take a system snapshot to retrieve the results.

Cluster Network Connectivity

- "Overview" on page 50
- "Configuring Cluster Network Connectivity Monitoring" on page 50

Overview

If you have a problem with a cluster, a Pulse Secure Support representative might ask you to enable the cluster node troubleshooting server. When you enable the server on the Network Connectivity tab, the system attempts to establish connectivity between the node on which the server resides and another node you specify. As the nodes communicate, the system displays network connectivity statistics on the page. The Network Connectivity tab is displayed only when you enable clustering on your system. On a standalone system, you do not have access to the Network Connectivity tab.

Use the Network Connectivity tab to enable the cluster node troubleshooting server and to select a node on which to perform troubleshooting tasks. The troubleshooting tool allows you to determine the network connectivity between cluster nodes.

The server component of this tool runs on the node to which connectivity is being tested. The client component runs on the node from which connectivity is being tested. The basic scenario for testing connectivity is this:

- The administrator starts the server component on the passive node.
- The administrator tests the connectivity to the server node from the Active node, by starting the client component on the active node and then contacting the passive node running the server component.

Note: The server component must be run on nodes that are configured as either standalone or in a cluster but disabled. Cluster services cannot be running on the same node as the server component.

Configuring Cluster Network Connectivity Monitoring

To enable network connectivity monitoring:

1. Select the **Enable cluster network troubleshooting server** check box to enable the server component.

Figure 30 shows the Troubleshooting page for network connectivity.

Figure 30 Troubleshooting using Network Connectivity

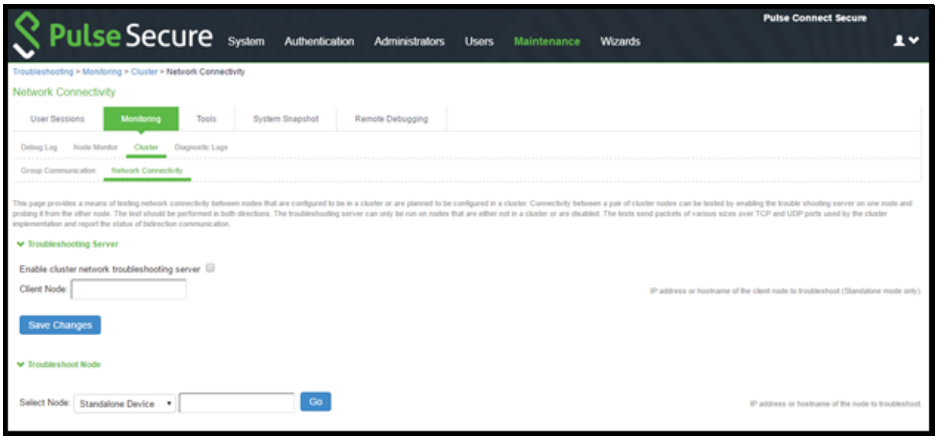

2. Click **Save Changes.**
3. On another machine, select **Maintenance > Troubleshooting > Cluster > Network Connectivity**.
4. Perform one of the following steps:

- Select a node from the list.
- Enter the IP address of the server node.

5. Click **Go** to begin troubleshooting the machine on which the server component is running.
6. Click the **Details** link below the fields to view the results.

WAN Clustering

Overview

A WAN cluster is a group of independent servers/nodes separated by WAN networks working together as a single system to provide load balancing and high scalability for clients and services. WAN cluster works only in active-active cluster operation mode, and is qualified on PSA7000, PSA7000-V, PSA5000, PSA5000-V and PSA3000 platforms.

Clustering supports following types of synchronization settings:

- **Configuration-only Cluster** - Only configuration will be synced across the cluster nodes
- **Synchronize user sessions** - Both configuration and user sessions will be synced across the cluster nodes

Note: WAN cluster only supports Configuration-only Cluster and does not support Synchronize user sessions.

Configuring an Active-Active Configuration-only WAN Cluster

To configure an active/active Configuration-only WAN Cluster:

1. First configure an active/active cluster as mentioned in the "Configuring an Active/Passive Cluster" on page 18 section.
2. Then, go to **System > Clustering > Cluster Properties** and select **Configuration-only Cluster** as shown in the screen below.

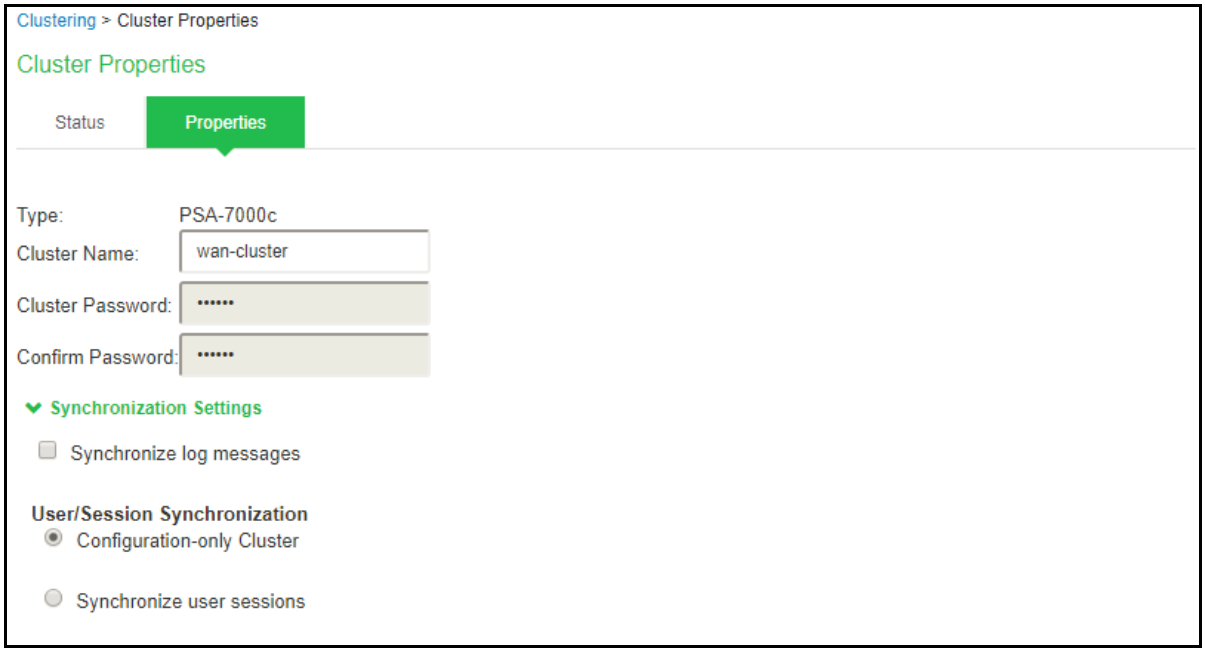

3. In the **Advanced Settings**, select the **Network Type** as **Average latency 60-100ms or Average latency 10-60ms for WAN cluster.** Refer to the image below.

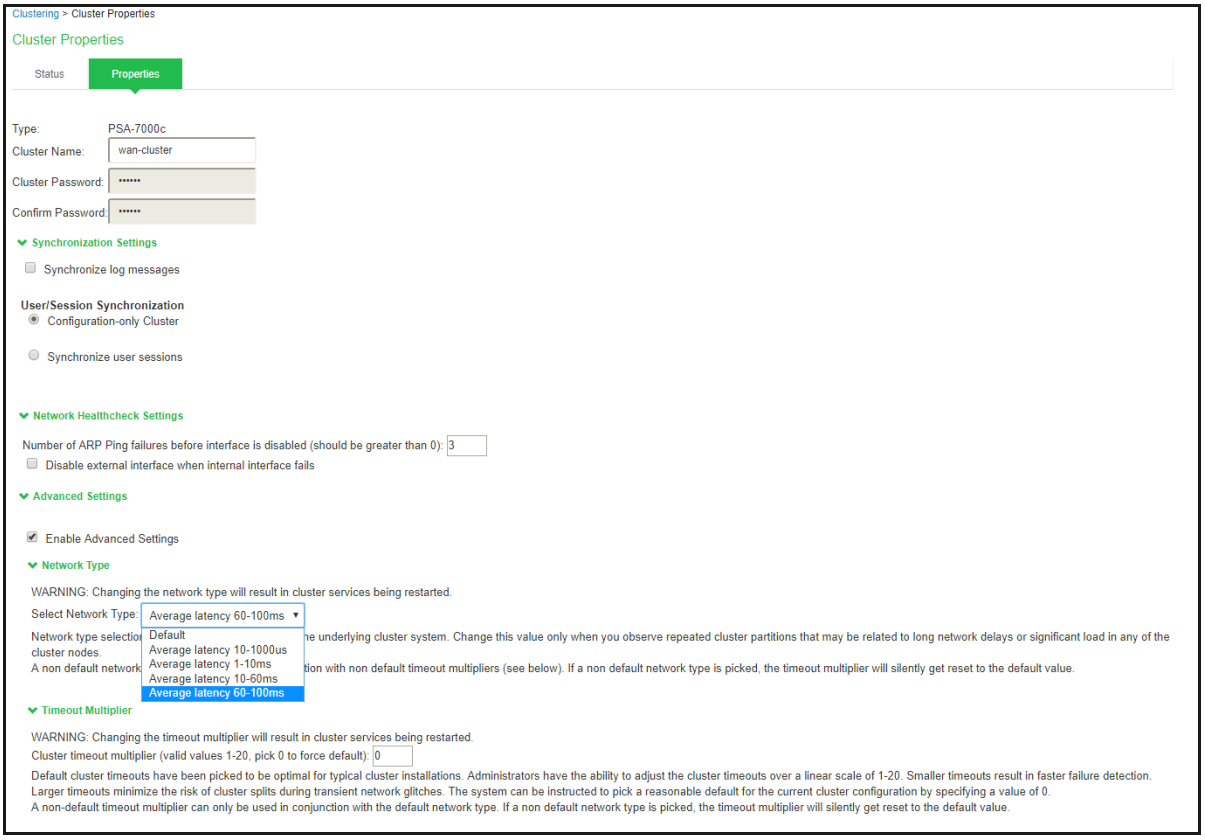

Note: For better performance a WAN cluster does not support configuring Global Static IP Pool VPN Connection Profile under Users -> Resource Policies -> VPN Tunneling -> Connection Profiles for Leasing IP to an end user client. Only Global DHCP IP Pool VPN Connection Profile Configuration or Node Specific Static/DHCP IP Pool VPN Connection Profile Configuration is supported.

Note: In an active/active WAN cluster, a connection profile configured with a Global Static IP Pool will be retained during Upgrade, Binary Import and XML Import with the below warning on the Dashboard and Overview Page for admin to take appropriate action. Also, an end user using VPN tunneling clients will not be leased IPs from the Global Static IP Pool.

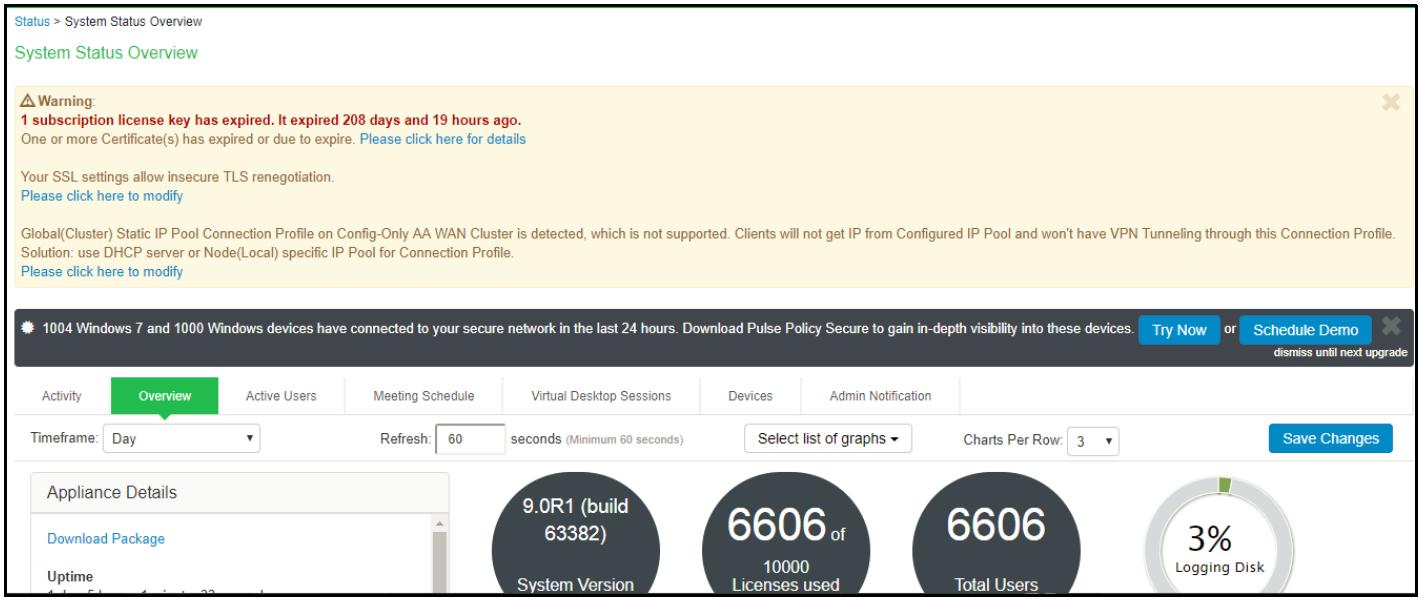

Note: In an active/active WAN cluster, if the networks of all the internal ports of the PCS/Nodes are in different subnets, it is mandatory to add specific static network routes on every PCS/Node to reach every other PCS/Node in the cluster for better cluster communication during PCS/Node failover or downtime.

To add a specific static route on a PCS/Node to reach another PCS/Node in the cluster:

1. Go to **System > Network > Routes.**

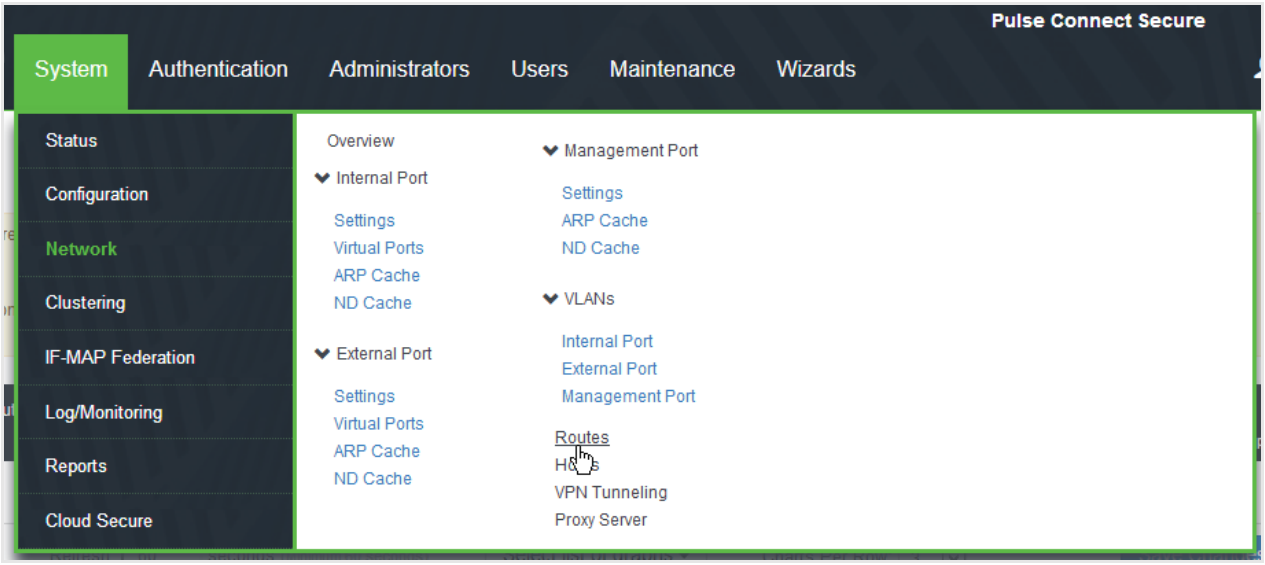

2. Click **New Route.**

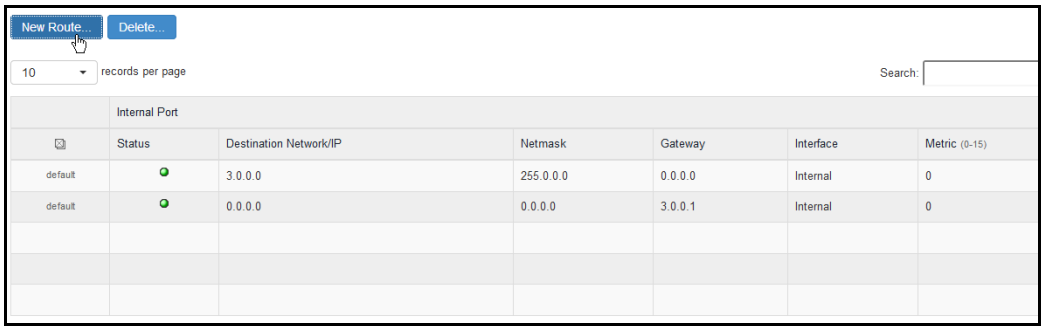

3. Based on the Network's Topology, the Static Route needs to be added on PCS/Node to reach other PCS/Node in WAN Cluster. Below is an example where static route is added on PCS Configured in 10.11.0.0/16 network having gateway 10.11.1.1 to reach another PCS/Node Configured in 10.12.0.0/16.

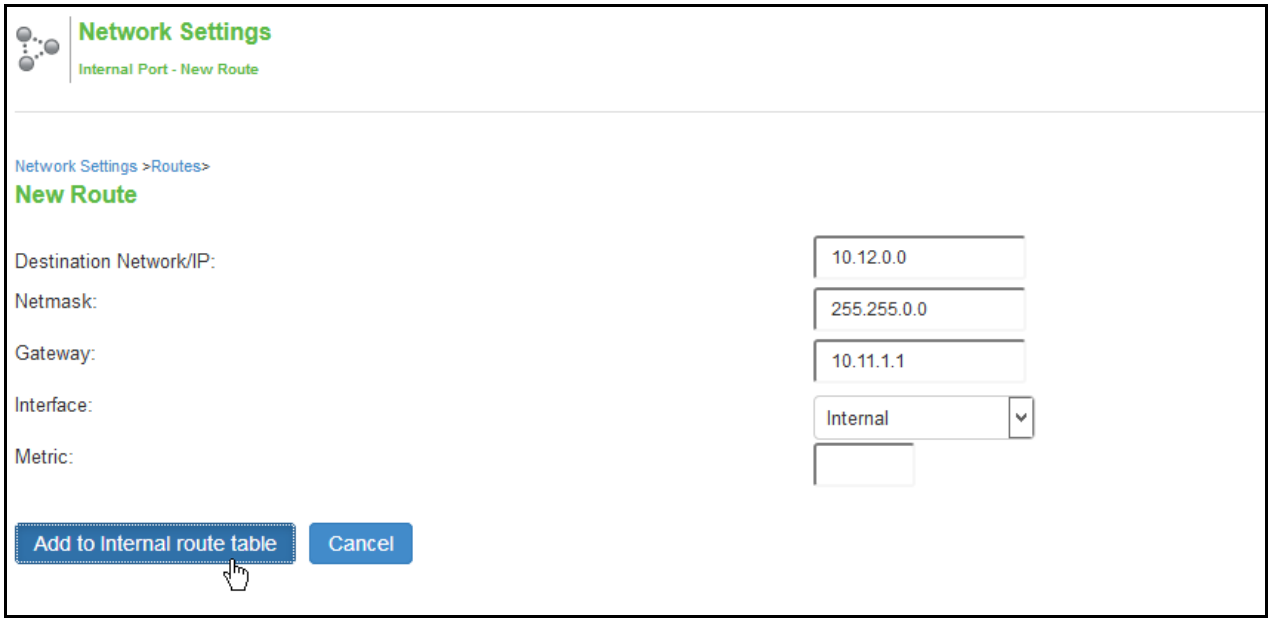

4. The same steps need to be repeated on every PCS/Node in the active/active WAN cluster.

Example: Creating an Active/Active Cluster That Supports IPv6 Client Access

This example describes the tasks involved in creating a cluster that supports IPv6 client access. It includes the following information:

Overview

Pulse Connect Secure supports an IPv6 configuration for active/active clusters. The previous intracluster communication mechanism is preserved. The intracluster communication occurs over the IPv4 corporate network through the internal interfaces.

If you attempt to change the IP address of a node while it belongs to a cluster, you might experience unpredictable results. Whenever you change the IP address configuration for a cluster, you must re-create the cluster. Therefore, to add support for IPv6 addresses, you must re-create the cluster.

Before You Begin

We recommend that you deploy a cluster in a staging environment first and then move to a production environment after testing authentication realm, user role, and resource policy configurations, as well as any applications your end users might access.

Before you begin a cluster configuration:

1. Ensure that all intended Pulse Connect Secure nodes use the same hardware platform (for example, all are PSA-7000C Appliances).
2. Ensure that all intended Pulse Connect Secure nodes have been initially configured (for example, Pulse Connect Secure hostname is specified, and the internal and external IP addresses are assigned), and they are running the same service package version.
3. Designate one node as the primary node. On the primary node, configure system and user settings. When other nodes join the cluster, the primary node propagates its configuration to the new cluster member during the join cluster operation.

Defining and Initializing a Cluster

You use the primary node admin GUI graphical user interface to create the cluster and add members. The primary node is added as part of the cluster creation operation. When you add members, you are prompted for settings unique to the member, such as the name and IP address configuration for the internal and external interfaces. A few additional settings are also unique, namely the management port and VLAN port settings, so you add these manually after the add node procedure that follows, but before the join cluster operation.

To create a cluster and add members:

1. Select **System > Clustering > Create** and enter a name for the cluster, a cluster password, and a name for this node, such as Node-1.

   You need to enter the password again when specifying additional nodes to join the cluster. All nodes in the cluster use this password to communicate.

2. Click **Create Cluster**. When prompted to confirm the cluster creation, click **Create**. After the Pulse Connect Secure initializes the cluster, the Clustering page displays the **Status** and **Properties** tabs.

3. Click **Add Members** to specify the additional cluster nodes:

   1. Enter a name for the member; for example, Node-2.

   2. Enter the internal IP address. If both IPv4 and IPv6 are enabled on the internal port on Node-1, the system prompts for both IPv4 and IPv6 settings for the internal port for Node-2. Note, however, that intracluster communication uses the IPv4 corporate network.

   3. Enter the external IP address. If both IPv4 and IPv6 are enabled on the external port on Node-1, the system prompts for both IPv4 and IPv6 settings for the external port for Node-2.

   4. Change the netmask/prefix-length and gateway settings for the node if necessary.

   5. Click **Add Node**. When prompted to confirm adding the new member, click **Add**.

      When the add node operation has completed, Node-2 is shown as an unreachable member of the cluster.

   6. The add node procedure does not prompt you to configure management port or VLAN port settings. As needed, go to the node port configuration page and configure these settings. For example, after the add node operation has completed for Node-2, go to its System > Network > Port > Settings page and configure its management port.

   7. Repeat this procedure for each node you intend to add to a cluster.

Joining Nodes to the Cluster

The primary node joins the cluster as part of the creation process. Use the following procedure to join additional nodes to the cluster.

To join a node to the cluster:

1. From an existing cluster member, select the **System > Clustering > Cluster Status** tab and specify the Pulse Connect Secure you want to add to the cluster.

2. From the admin GUI of the Pulse Connect Secure you want to join to a cluster:

1. Select the **System > Clustering > Join** tab and enter:
   - The name of the cluster to join.
   - The cluster password you spehscified when defining the cluster.
   - The IPv4 address for the internal port of an active cluster member.
2. Click **Join Cluster.** When prompted to confirm joining the cluster, click **Join**.

The join cluster operation validates IPv4 and IPv6 settings for all the physical ports (internal/external/management) against those present in the existing cluster. For example, the external port IPv6 settings present on Node-2 are compared against external port IPv6 settings that were specified for the Node-2 add member operation entered on the primary node (Node-1). If there is a mismatch, the join operation fails with an appropriate error message.

While the new node synchronizes its state with the existing cluster member, each node's status indicates **Enabled, Enabled, Transitioning, or Enabled, Unreachable.**

When the node finishes joining the cluster, its Clustering page shows the **Status** and **Properties** tabs.

After the node joins the cluster, you might need to sign in again.

Advanced Configuration

Table 8 summarizes advanced configuration guidelines.

Table 8 Pulse Connect Secure Clusters: Advanced Configuration Guidelines

Topic	Guideline
Active/Active	When using Pulse Secure clients with an active/active cluster, you must split the IP address pool across the nodes to ensure proper routing from the backend to the end user. This is a requirement whether the IP address pool is provisioned statically on the Pulse Connect Secure or dynamically by way of DHCP. The client IP pool configuration is synchronized among all nodes in a cluster; however, you may configure each node to use a certain subset of the global IP pool. If you are running Network Connect on a multisite cluster where nodes reside on different subnets: 1 Configure an IP address pool policy on the Users > Resource Policies > VPN Tunneling: Connection Profiles > New Profile page that accounts for the different network addresses used by each node in the cluster. 2 For each node in the cluster, use settings in the System > Network > VPN Tunneling page of the admin GUI to specify an IP filter that filters out only those network addresses available to that node. 3 Create a static route on your gateway router that indicates the IP address of the internal port of each cluster node. Each IP address specified on the router needs to be in the same subnetwork as the corresponding cluster node.
FIPS	If you are creating a cluster of FIPS devices, manually update the security world on each of the nodes.

Example: Creating an Active/Passive Cluster that Supports IPv6 Client Access

This example describes the tasks involved in creating a cluster that supports IPv6 client access. It includes the following information:

Overview

Pulse Secure access management framework supports an IPv6 configuration for active/passive clusters. The previous intracluster communication mechanism is preserved. The intracluster communication occurs over the IPv4 corporate network through the internal interfaces.

If a device belongs to an active/passive cluster, you can enable IPv6 on its ports. If a device has IPv6 enabled on its ports, it can be added to an active/passive cluster.

If you attempt to change the IP address of a node while it belongs to a cluster, you might experience unpredictable results. Whenever you change the IP address configuration for a cluster, you must re-create the cluster.

When using active/passive clustering, the members of a cluster pair must be in the same subnet because the VIP address must be shared by both members.

Before You Begin

We recommend that you deploy a cluster in a staging environment first and then move to a production environment after testing the authentication realm, user role, and resource policy configurations, as well as any applications your end users might access.

Before you begin a cluster configuration:

Note that state synchronization occurs only through the internal network interface card (NIC).

Ensure that all intended nodes use the same hardware platform (for example, all are PSA-7000C Appliances).

Ensure that all intended nodes have been initially configured (for example, the system hostname is specified, and the internal and external IP addresses are assigned), and that they are running the same service package version.

Designate one node as the primary node. On the primary node, configure system and user settings. When other nodes join the cluster, the primary node propagates its configuration to the new cluster member during the join cluster operation.

Configuring IPv6 on an existing IPv4 active/passive cluster on an external port can be done seamlessly. However, if you are configuring on an internal port, you must wait for cluster synchronization completion and then do the next configuration for the remaining node. Therefore, we recommended that you complete the IPv6 configurations before creating a cluster on an internal port.

Defining and Initializing a Cluster

You use the primary node admin GUI to create the cluster and add members. The primary node is added as part of the cluster creation operation. When you add members, you are prompted for settings unique to the member, such as the name and IP address configuration for the internal and external interfaces. A few additional settings are also unique, namely the management port and VLAN port settings, so you add these manually after the add node procedure that follows, but before the join cluster operation.

To create a cluster and add members:

1. Select **System > Clustering > Create Cluster** and enter a name for the cluster, a cluster password, and a name for this node, such as Node-X.

   You need to enter the password again when specifying additional nodes to join the cluster. All nodes in the cluster use this password to communicate.

Figure 31 shows the Create New Cluster page.

Figure 31 Create New Cluster Page

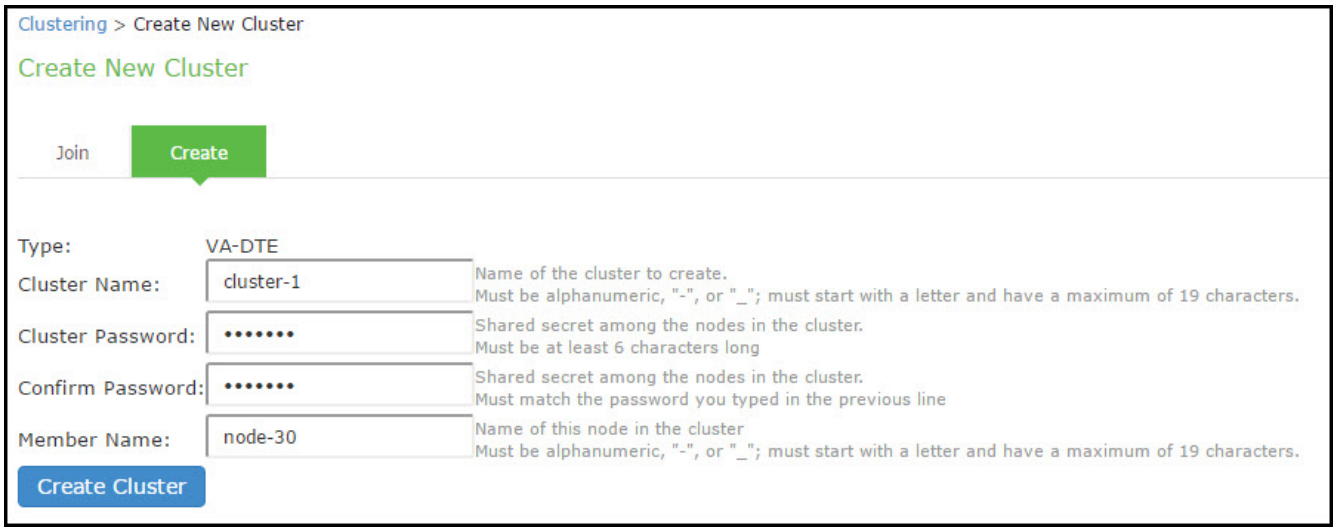

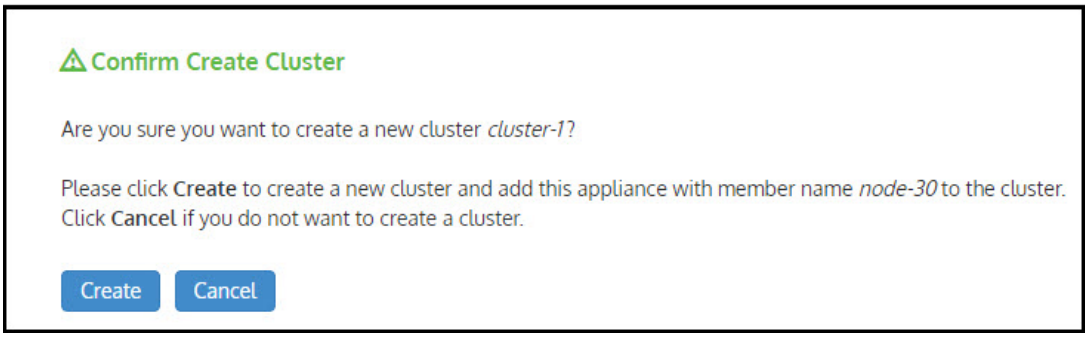

2. Click **Create Cluster.** When prompted to confirm the cluster creation, click **Create**. After Connect Secure initializes the cluster, the Clustering page displays the **Status** and **Properties** tabs.

Figure 32 shows the Clustering page with Status and Properties tabs.

Figure 32 Clustering Page- Status and Properties

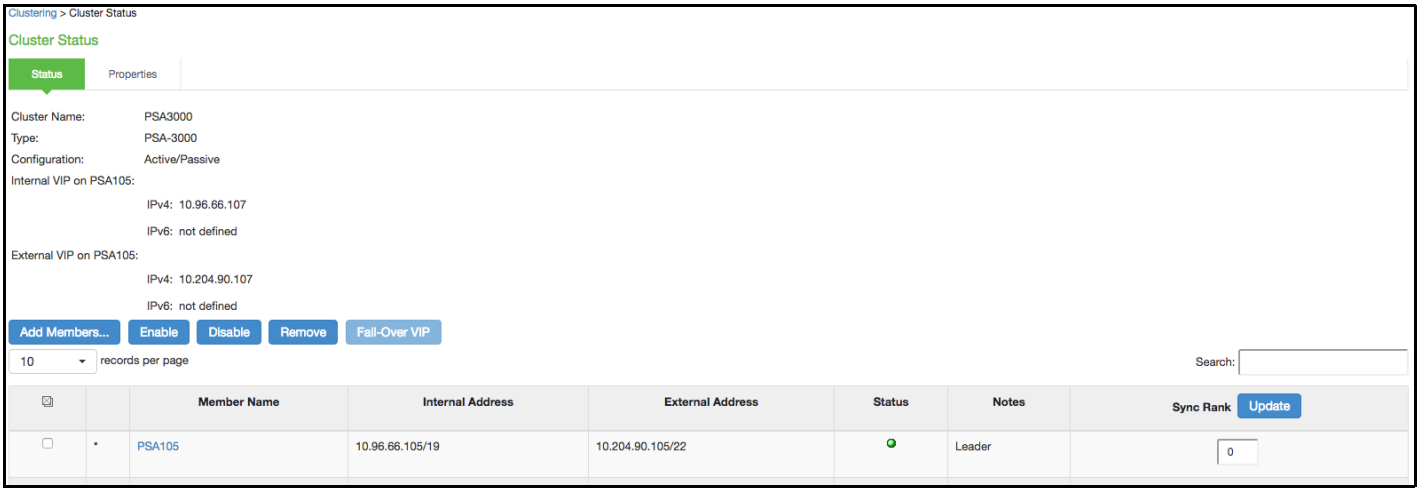

3. Click **Properties**.

Figure 33 shows the Clustering page with active/passive configuration.

Figure 33 Clustering Page- Active/Passive Configuration

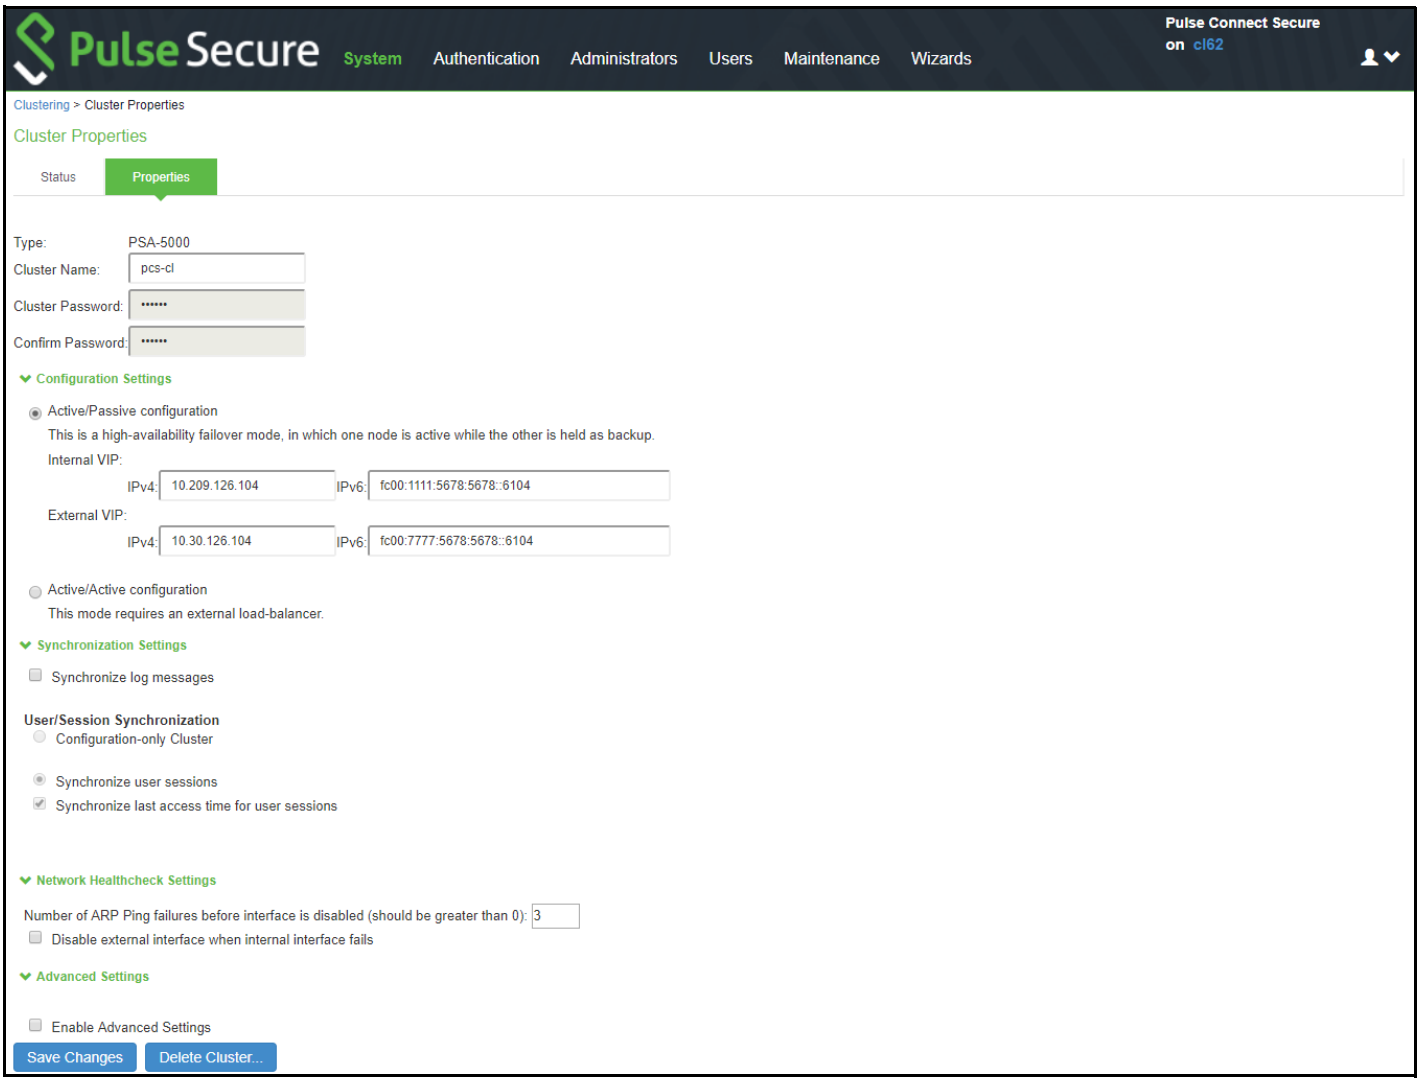

4. Under Configuration Settings, select **Active/Passive Configuration,** then specify the IPv4 and IPv6 addresses for the VIP address on the internal and external ports, depending on what is enabled for **IPv4/IPv6 at Network > Internal Port and Network > External Port.**

5. Click **Save Changes.** After the system initializes the active/passive cluster, the Clustering page displays the **Status** and **Properties** tabs.

6. Click **Add Members** to specify additional cluster nodes:

Figure 34 Add Cluster Member Page

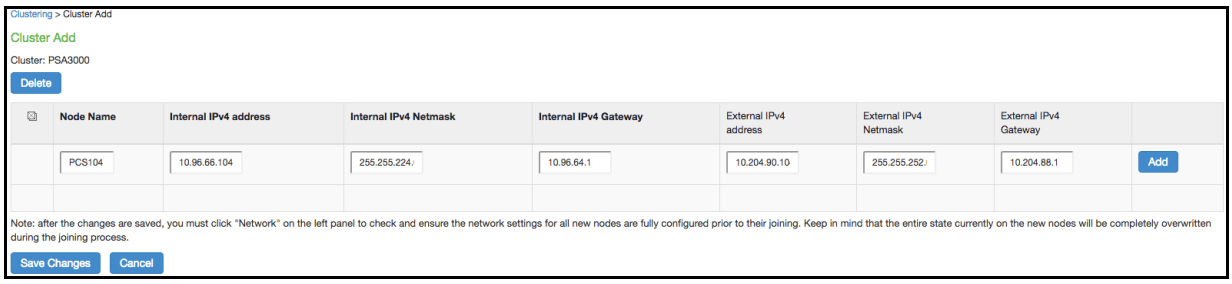

7. Enter a name for the member; for example, Node-Y.
8. Enter the internal IP address. If both IPv4 and IPv6 are enabled on the internal port on Node-X, the system prompts for both IPv4 and IPv6 settings for the internal port for Node-X. Note, however, that intracluster communication uses the IPv4 corporate network.
9. Enter the external IP address. If both IPv4 and IPv6 are enabled on the external port on Node-X, the system prompts for both IPv4 and IPv6 settings for the external port for Node-Y.
10. (Optional) Change the netmask, prefix-length, and gateway settings for the node if necessary.
11. Click **Add Node**. When prompted to confirm adding the new member, click Add and then click **Save Changes**.
12. After the completion of add node operation, Node-Y is shown as an unreachable member of the cluster.
13. Verify the configuration on **System > Clustering > Cluster Status** page.

Figure 35 shows the status on the Clustering page.

Figure 35 Clustering Page -Status

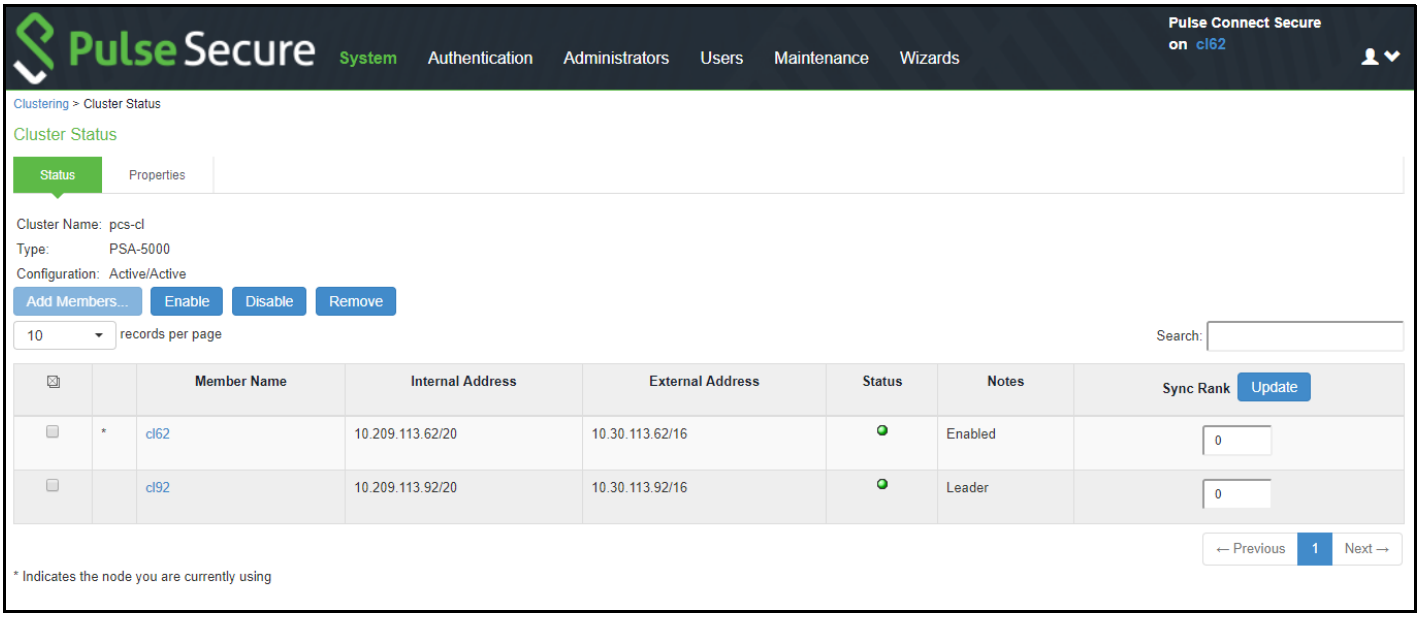

The add node procedure does not prompt you to configure management port or VLAN port settings. As needed, go to the node port configuration page and configure these settings. For example, after the add node operation has completed for Node-Y, go to its System > Network > Port > Settings page and configure its management port.

Note: Only two nodes can be present in an active/passive cluster.

Joining Nodes to the Cluster

The primary node joins the cluster as part of the creation process.

To join additional nodes to the cluster:

1. From an existing cluster member, select the **System > Clustering > Cluster Status** tab and specify the Connect Secure you want to add to the cluster.
2. From the admin GUI of the Pulse Secure access management framework that you want to join to a cluster:
    1. Select the **System > Clustering > Join** tab and enter:
        - The name of the cluster to join
        - The cluster password you specified when defining the cluster
        - The IPv4 address for the internal port of an active cluster member
    2. Click **Join Cluster**. When prompted to confirm joining the cluster, click **Join**.

The join cluster operation validates IPv4 and IPv6 settings for all the physical ports (internal, external, and management) against those present in the existing cluster. For example, the external port IPv6 settings present on Node-Y are compared against external port IPv6 settings that were specified for the Node-Y add member operation entered on the primary node (Node-X). If there is a mismatch, the join operation fails with an appropriate error message.

While the new node synchronizes its state with the existing cluster member, each node's status indicates Enabled, Enabled, Transitioning, or Enabled, Unreachable.

When the node finishes joining the cluster, its Clustering page shows the **Status** and **Properties** tabs.

After the node joins the cluster, you might need to sign in again.

Configuring IPv6 on an Existing IPv4 Active/Passive Cluster

We recommend as a best practice that you configure IPv6 host and network settings on individual nodes before you create a cluster. In some cases, such as routine upgrade, you have already created a cluster configuration and only want to add IPv6 addresses to the existing interface configuration. If so, follow the procedures in this section precisely.

Note: You must leave IPv6 disabled until the last step of the procedures shown below.

To modify the internal port configuration for the cluster:

1. Select **System > Network > Internal Port > Settings.**
2. Under Settings for, select **Entire cluster**.
3. Complete the configuration for the IPv6 prefix and the IPv6 gateway, but do not enable IPv6.
4. Verify that all the nodes are up and running, are in sync, and are in reachable state. Complete synchronization of the cluster pair might take a few minutes.
5. Under Settings for, select **Node 1.**
6. Configure the IPv6 address, but do not enable IPv6.
7. Verify both the nodes are up and running and in reachable state.
8. Repeat steps 6-8 for Node 2.
9. Select **System > Network > Internal Port > Virtual Ports.**
10. Update the cluster virtual port configuration to add the IPv6 address.
11. Select **System > Network > Internal Port > Settings.**
12. Under Settings for, select **Entire cluster.**
13. Select **Enable IPv6.**

To modify the external port configuration for the cluster:

1. Select **System > Network > External Port > Settings.**
2. Under Settings for, select **Entire cluster**.
3. Complete the configuration for the IPv6 prefix and the IPv6 gateway, but do not enable IPv6.
4. Verify that all the nodes are up and running, are in sync, and are in reachable state. Complete synchronization of the cluster pair might take a few minutes.
5. Under Settings for, select **Node 1**.
6. Configure the IPv6 address, but do not enable IPv6.
7. Verify both the nodes are up and running and in reachable state.
8. Repeat steps 6-8 for Node 2.
9. Select **System > Network > External Port > Virtual Ports.**
10. Update the cluster virtual port configuration to add the IPv6 address.
11. Select **System > Network > External Port > Settings.**
12. Under Settings for, select **Entire cluster.**
13. Select **Enable IPv6.**

Advanced Configuration

Table 9 summarizes advanced configuration guidelines.

Table 9 Pulse Connect Secure Clusters: Advanced Configuration Guidelines

Settings	Guideline
FIPS	If you are creating a cluster of FIPS devices, manually update the security word on each of the nodes.