

How to deploy IVE Active-Active and Active-Passive clusters

Overview

Juniper Netscreen SA and SM series appliances support Active/Passive or Active/Active configurations across a LAN or a WAN to provide high availability, increased scalability, and load balancing capabilities.

You define a cluster on one IVE by configuring:

- A name for the cluster
- A password for the cluster members to share
- A name to identify the machine in the cluster

After specifying this information on the **System > Clustering > Create** tab, you click **Create Cluster** to initiate the cluster and add the current machine to the cluster. After creating the cluster, the **Clustering** page shows **Status** and **Properties** tabs, which replace the original, **Join** and **Create** tabs.

The **Status** tab lists the cluster name, type, and configuration (active/active or active/passive), enables you to specify new members and manage existing members, and provides overall cluster status information. The **Properties** tab enables you to change the cluster name and set configuration, synchronization, and health-check settings.

After defining and initializing a cluster, you need to specify which IVEs will be added to the cluster. After an IVE is identified as an intended member, you may add it to the cluster through the following modes.

1. Web console—If a configured IVE is running as a stand-alone machine, you can add it to a cluster through its Web console.
2. Serial console—If an IVE is in its factory-state, you can add it to a cluster through its serial console by entering minimal information during initial setup.

When an IVE joins a cluster, it initializes its state from the existing member that you specify. The new member sends a message to the existing member requesting synchronization. The existing member sends the system state to the new member, overwriting all system data on that machine. After that point, the cluster members synchronize data when there is a state change on any member.

Cluster member communication is encrypted to prevent attacks from inside the corporate firewall. Each IVE uses the shared password to decrypt communication from another cluster member. For security reasons, the cluster password is not synchronized across IVEs.

Note: During synchronization, the new node receives the service package, which upgrades the node if it is equipped with a Central Manager license and is running an older service package.

Configuration

Create an IVE cluster:

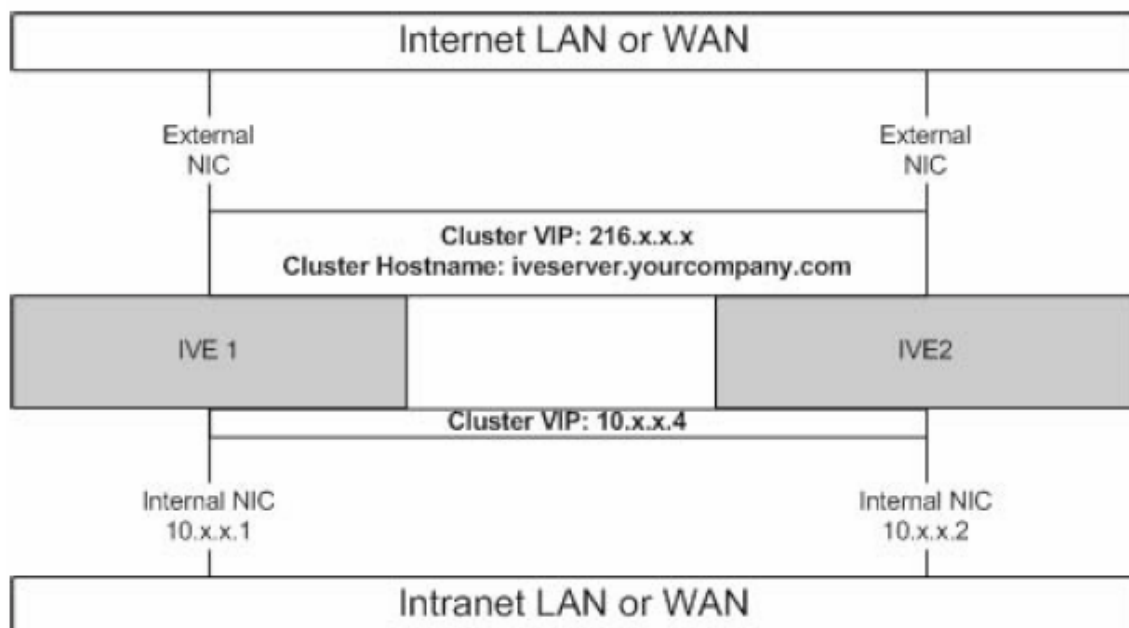
1. Initialize the IVE cluster through the **System > Clustering > Create Cluster** page of the Web Console by defining the cluster name and adding the first/primary IVE to the cluster.
2. Add the names and IP addresses of future cluster IVEs to the primary IVE through the **System > Clustering > Status** page of the Web Console.
3. Populate the cluster with additional IVEs as necessary through the **System > Clustering > Join Cluster** page of the Web Console.

Deploying two nodes in an Active/Passive cluster

You can deploy IVEs as a cluster pair in Active/Passive mode. In this mode, one IVE actively serves user requests while the other IVE runs passively in the background to synchronize state data, including system state, user profile, and log messages.

User requests to the cluster VIP (virtual IP address) are passed to the active IVE. If the active IVE goes off-line, the standby IVE automatically starts servicing user requests. Users do not need to sign in again, however some IVE 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 IVE box, in which case users may need to sign in to back-end Web servers again.

The following diagram illustrates an Active/Passive IVE cluster configuration using two IVEs that have enabled external ports. Note that this mode does not increase throughput or user capacity, but provides redundancy to handle unexpected system failure.



Active/Passive Cluster Pair

Deploying two or more units in an Active/Active cluster

In Active/Active mode, all the machines in the cluster actively handle user requests sent by an external load balancer or Round-Robin DNS. The load balancer hosts the cluster VIP and routes user requests to an IVE defined in its cluster group based on source-IP routing.

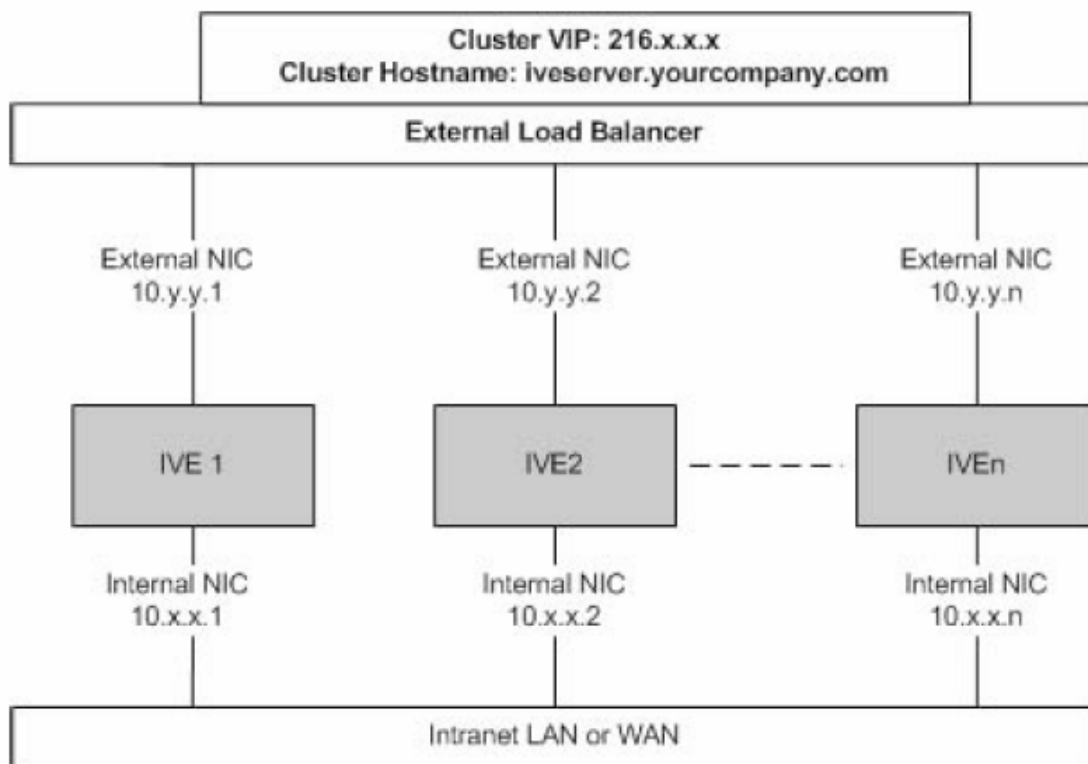
If an IVE goes off-line, the load balancer adjusts the load on the active IVEs. Users do not need to sign in again, however some IVE session information entered a few seconds before the active machine went off-line, such as cookies and passwords, may not have been synchronized on the current IVE box, in which case users may need to sign in to back-end Web servers again.

The IVE 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 the cluster members. When an off-line IVE 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 IVE hosts an HTML page that provides service status for each IVE in a cluster. External load balancers can check this resource to determine how to effectively distribute the load among all the cluster nodes.

To perform the L7 Health Check for a node:

- **From a browser**—Enter the following URL:
https://<IVE-Hostname>/dana-na/healthcheck/healthcheck.cgi
- **Using an external load balancer**—Configure a Health Check policy that sends the following request to cluster nodes:
GET /dana-na/healthcheck/healthcheck.cgi HTTP/1.1\nHost: localhost
The node returns one of two values:
 1. **“Cluster Enabled” string**—this value means that the node is active.
 2. **“500”**—this value denotes an error and cluster IVEs stop forwarding user requests to the node.



Active/Active IVE cluster configuration in which the IVEs have enabled external ports.

State Synchronization:

IVE state synchronization occurs only via the internal network interface cards (NICs), and each cluster member is required to possess the cluster password in order to communicate with other members. Cluster members synchronize data when there is a state change on any member. IVE cluster state data is either *'persistent'*—permanently stored on the IVE—or *'transient'*—stored on the IVE only for the user's session.

IVE state data is divided into the following major categories:

1. **System state**—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
2. **User Profile** - This data (bookmarks, persistent user cookies and persistent user passwords) can be either persistent or transient, depending on whether or not you have enabled persistent cookies and persistent password caching.
 3. **User Session** - This state is transient and dynamic (IVE session cookie and other user profile information which is stored only for a session).
 4. **Monitoring state**—Persistent information consists of log messages. Please note that when you add an IVE 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 off-line machine comes back online. Once all machines are online, however, log messages are synchronized.

Deploying a cluster in an Access Series FIPS environment

In addition to sharing state, user profile, user session, and monitoring state data, the members of an Access Series FIPS cluster also share security world data.

All cluster members share the same private key and are accessible using the same administrator cards. Since changing a security world requires physical access to a cryptographic module, however, Access Series FIPS cluster members cannot share all of their data using the standard IVE synchronization process. Instead, to create an Access Series FIPS cluster, you must:

1. *Create a cluster of Access Series FIPS machines through the Web console—*

As with a standard IVE cluster, each cluster node in an Access Series FIPS cluster is initialized using system state data from the specified cluster member, overwriting all existing data on the node machine.

2. *Manually update the security world on each of the machines—*after creating a cluster, you must initialize each cluster node with the specified member's security world using an administrator card that is pre-initialized to the security world, a smart card reader, and the serial console.

Similarly, if you want to modify an existing security world on a cluster, you must individually update each cluster member's cryptographic module using an administrator card, smart card reader, and the IVE serial console.

FAQs

1. A node up and running appears in the WEB UI as unreachable
Answer: A cluster member may appear as unreachable even when it is online and can be pinged.
Here are reasons why a node can show as unreachable:
 - its password is incorrect. If the node has never joined the cluster or if the password has changed in between this might be a possibility
 - it does not know about all the nodes of the cluster
 - it has different group communication mode
 - it has a different version
2. Both my machines in the cluster are up and in a cluster but each indicate that the other is unreachable.
Answer: Check item #1. While the clusters seem the same they do have something different, e.g. password, version etc
3. After I join the machine to a running cluster my session times out and I have to login again
Answer: This is expected behavior. The member that joins the cluster get all its state (including the active sessions) overwritten by the state in the cluster. Therefore the session you use to join the cluster is closed.

4. I created a cluster and added a node but it appears as Unreachable. When I login to the other machine it does not appear to be member of the cluster
Answer: It is not enough to just add a machine in the cluster in a machine already in the cluster. You also need to go to the machine that is being added to the cluster and use the join UI to add the machine to the cluster.
 The IVE provides such a UI in two places: a) Part of the WEB UI b) part of the console UI when a machine boots
5. What protocols and ports are used for clustering?

Protocol	Port	When	Purpose
TCP/IP	4808	Clustering on, Always	P2P encrypted communication
	4809	Clustering on, Always	P2P clear text communication
	4900-4910	For a short period during handshake	Key exchange for group communication, state sync where applicable
UDP	4803	Clustering On, always	Group communication
	4804	Clustering On, always	Token Heartbeat

6. In the cluster status page, when I hover the mouse over the status gif I see a hexadecimal number. What is the meaning of this number?

Answer: The hexadecimal number is a snapshot of the status of the IVE. It is a bit mask indicating a number of states as shown in the table below. Each bit in the bit mask represents a sub state. Good state – 0x18004 on one node and 0x10004 on rest.

Value	Meaning
0x000001	IVE in standalone mode
0x000002	IVE in cluster disabled state
0x000004	IVE in cluster enabled state
0x000008	IVE is unreachable (because it is offline, broken network connectivity, password mismatch , has different cluster definition, software version mismatch etc)
0x000100	IVE is syncing state from another IVE (initial syncing phase)
0x000200	IVE is transitioning from one state to another
0x000800	IVE eth0 appears disconnected (no carrier)
0x001000	IVE eth1 appears disconnected (no carrier)
0x002000	IVE is syncing its state to another IVE that is joining
0x004000	Initial Synchronization as master or slave is going on
0x008000	This IVE is the leader of the cluster
0x010000	The spread daemon is running and the cache server is connected to it
0x020000	The gateway on eth0 is unreachable for ARP pings (see log file)
0x040000	The gateway on eth2 is unreachahble for ARP pings (see log file)
0x800000	Leader Election is taking place
0x100000	Server Lifecycle process (dsmond) is busy
0x200000	System is performing post state synchronization activities

7. For all incoming https requests answered by the external VIP, what would be the source IP and Mac Address for the reply packets?
Source IP will be the VIP Address and the MAC address will be that of the active IVE that responds.
8. How does log synchronization work in an IVE A/P cluster?
If log sync is enabled, only the logserver leader (not necessarily the active node) will send logs to the syslog. The non-leader will send its messages to the leader, thereby ensuring all messages end up at the syslog server.
If sync is disabled, both the nodes will independently send their messages to the syslog server.

9. How does the IVE decide the Logserver leader and is there a way to identify the leader at a specific time?

Cluster members use a number of heuristics to pick the cluster leader. The goal is to attempt to designate a node that is expected to have the most recent state about the cluster as the leader.

For the admin the easiest way to determine who the leader of a cluster at any point of time is to go to the System->Clustering->Status page and hover the mouse over the bullets under the "Status" column. When the mouse hovers over a bullet, the system will show a hexadecimal number for each bullet. The node that has the 0x8000 bit on is the leader. There will be only one node with this bit turned ON. In other words, the fourth least significant digit in the hexadecimal number for any node is any of 8, 9, a, b, c, d or e, then the node is the leader.

Notes:

- **Upgrading clusters:**
 1. With Central Manager (Central Manager is a licensable feature) – Central Manager will detect the upgrade of a single node in the cluster, and upon its reboot/re-synch, it will instruct the other nodes to upgrade themselves automatically by sending them the service package.
 2. Without Central Manager – To upgrade nodes in a cluster, the Admin should disable the clustered nodes, upgrade each node individually, and after the nodes reboot, re-enable them in the cluster.
- **Restarting or rebooting clustered nodes:**
 - When you create a cluster of two or more IVEs, the clustered IVEs acts as a logical entity. As such, when you restart or reboot one of the clustered IVEs using either the serial console or the Web console, all IVEs in the cluster restart or reboot.
 - If you want to restart or reboot only one IVE in a cluster, first use the controls on the System > Clustering > Status page to disable the IVE you want to restart or reboot within the cluster. Next, use the controls on the Maintenance > System > Platform page, or the serial console's Reboot/Shutdown/Restart this IVE menu item, to restart or reboot the IVE. After the IVE restarts or reboots, enable the IVE within the cluster again.