Citrix NetScaler Clustering Guide

Citrix® NetScaler® 10

Copyright and Trademark Notice

ALTHOUGH THE MATERIAL PRESENTED IN THIS DOCUMENT IS BELIEVED TO BE ACCURATE, IT IS PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE ALL RESPONSIBILITY FOR THE USE OR APPLICATION OF THE PRODUCT(S) DESCRIBED IN THIS MANUAL.

CITRIX SYSTEMS, INC. OR ITS SUPPLIERS DO NOT ASSUME ANY LIABILITY THAT MAY OCCUR DUE TO THE USE OR APPLICATION OF THE PRODUCT(S) DESCRIBED IN THIS DOCUMENT. INFORMATION IN THIS DOCUMENT IS SUBJECT TO CHANGE WITHOUT NOTICE. COMPANIES, NAMES, AND DATA USED IN EXAMPLES ARE FICTITIOUS UNLESS OTHERWISE NOTED.

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

Modifying the equipment without Citrix' written authorization may result in the equipment no longer complying with FCC requirements for Class A digital devices. In that event, your right to use the equipment may be limited by FCC regulations, and you may be required to correct any interference to radio or television communications at your own expense.

You can determine whether your equipment is causing interference by turning it off. If the interference stops, it was probably caused by the NetScaler appliance. If the NetScaler equipment causes interference, try to correct the interference by using one or more of the following measures:

Move the NetScaler equipment to one side or the other of your equipment.

Move the NetScaler equipment farther away from your equipment.

Plug the NetScaler equipment into an outlet on a different circuit from your equipment. (Make sure the NetScaler equipment and your equipment are on circuits controlled by different circuit breakers or fuses.)

Modifications to this product not authorized by Citrix Systems, Inc., could void the FCC approval and negate your authority to operate the product.

BroadCom is a registered trademark of BroadCom Corporation. Fast Ramp, NetScaler, and NetScaler Request Switch are trademarks of Citrix Systems, Inc. Linux is a registered trademark of Linus Torvalds. Internet Explorer, Microsoft, PowerPoint, Windows and Windows product names such as Windows NT are trademarks or registered trademarks of the Microsoft Corporation. NetScape is a registered trademark of Netscape Communications Corporation. Red Hat is a trademark of Red Hat, Inc. Sun and Sun Microsystems are registered trademarks of Sun Microsystems, Inc. Other brand and product names may be registered trademarks or trademarks of their respective holders.

Portions of this software may be redistributed under an open source license. Information about those portions of the software, including a listing of all third party attribution notices and open source license agreements can be found at http://www.citrix.com/lang/English/lp/lp_2305124.asp.

Last Updated: July 2012

Document code: July 19 2012 02:03:50

Contents

Clustering

A NetScaler cluster is a group of NetScaler nCore appliances working together as a single system image. Each appliance of the cluster is called a *node*. A NetScaler cluster can include as few as 2 or as many as 32 NetScaler nCore hardware or virtual appliances as nodes.

The client traffic is distributed between the nodes to provide high availability, high throughput, and scalability.

To create a cluster, you add the required NetScaler appliances as cluster nodes, set up communication between the nodes, set up links to the client and server networks, configure the NetScaler appliances, and configure the distribution of client and server traffic.

NetScaler Features Supported by a Cluster

The following table lists the NetScaler features that are fully supported on the cluster, features that only work on individual cluster nodes, and features that are not supported on the cluster.

Table 1-1. NetScaler feature support matrix

Supported features	Features supported at node-level	Unsupported features
◆ Load Balancing ◆ Load Balancing Persistency ◆ SIP ◆ maxClient ◆ Spillover ◆ SSL PI policy ◆ Content Switching ◆ Cache Redirection ◆ Compression Control ◆ Content Filtering ◆ OSPF (IPv4 and IPv6) ◆ RIP (IPv4 and IPv6) ◆ BGP (IPv4 and IPv6)	◆ Surge Protection ◆ Sure Connect ◆ Priority Queuing ◆ HTTP Denial-of-Service Protection (HTTP DoSP) ◆ Integrated Caching ◆ Call Home	◆ DNS Load Balancing ◆ FTP Load Balancing ◆ Global Server Load Balancing (GSLB) ◆ NetScaler Push ◆ RTSP ◆ Stateful Connection Failover ◆ Graceful Shutdown ◆ DBS Auto Scaling ◆ DSR using TOS ◆ Spillover based on bandwidth ◆ Finer Startup-RR Control

Supported features	Features supported at node-level	Unsupported features
◆ HTML Injection ◆ TCP Buffering ◆ Distributed Denial-of-Service (DDoS) ◆ Basic Networking (IPv4 and IPv6) ◆ VLAN ◆ ICMP ◆ Fragmentation ◆ MAC-Based Forwarding (MBF) ◆ RNAT ◆ INAT ◆ KRPC ◆ ACL ◆ Simple ACL ◆ PBR ◆ SNMP GET/SET, Walk ◆ SNMP Traps ◆ Policy Infrastructure (PE/PI) ◆ NITRO API ◆ AppExpert ◆ Rewrite ◆ Responder ◆ Use Source IP (USIP) ◆ AppFlow exporter and Appflow collector (client) with waterfall chart ◆ DataStream ◆ MSR		◆ Rate Limiting ◆ Stream Analytics ◆ Net Profile ◆ DNS Caching ◆ SSL-VPN ◆ SSL CPE policy ◆ Application Firewall ◆ AAA ◆ Cloud Bridging - Tunneling ◆ Layer2 Mode ◆ FIPS ◆ XML XSM ◆ AAA-TM ◆ VMAC/VRRP ◆ Link Load Balancing ◆ IP-IP Tunneling ◆ DHCP RA ◆ Bridge Group ◆ Network Bridge ◆ Web Interface on NetScaler ◆ EdgeSight Monitoring ◆ BR LB ◆ ISIS Routing ◆ FIS (Failover Interface Set)

Supported features	Features supported at node-level	Unsupported features
- Policy-based RNAT - Web logging - Auditing (syslog and nsauditlog) - Path MTU Discovery - Client Keep-alive		

Hardware and Software Requirements

Appliances that you add to a NetScaler cluster must satisfy the following requirements:

- Be NetScaler nCore appliances. Clustering of NetScaler Classic appliances is not supported.
- Be of the same platform type (physical appliance or VPX instance).
- Be of the same hardware type (for physical appliances).
- Be on the same subnet.
- Have the cluster license file.
- Have the same licenses (Standard, Enterprise, or Platinum, and any add-on licenses).
- Be of the same software version and build.
- Be initially configured and connected to a common client-side and server-side network.

How Clustering Works

A NetScaler cluster is formed by grouping NetScaler appliances that satisfy requirements specified in Hardware and Software Requirements on page 9. One of the cluster nodes is designated as a *configuration coordinator (CCO)*. As the name suggests, the CCO coordinates all cluster configurations through the management IP address of the cluster, which is called the *cluster IP address*.

The cluster must be configured by accessing the CCO through the cluster IP address, as shown in the following figure:

Figure 1-1. Configuring the cluster through the cluster IP address

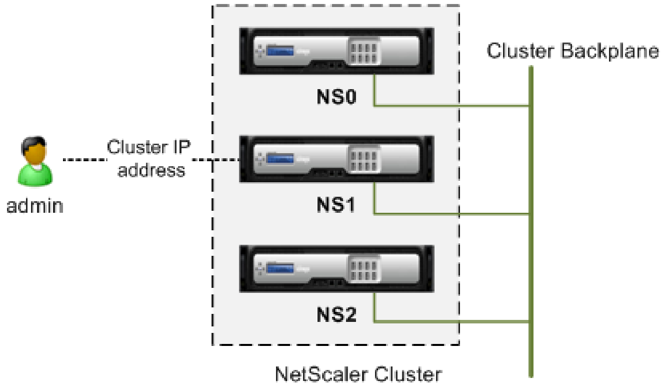

Note: You cannot configure an individual node by accessing it through the NetScaler IP (NSIP) address. Nodes accessed through the NSIP address are available in read-only mode. This means that you can only view the configurations and the statistics. However, there are some commands that can be executed on individual nodes. For more information, see Operations Supported on Individual Nodes on page 63.

The VIP addresses that you define on a cluster are available on all the nodes of the cluster (*striped addresses*). You can define SNIP addresses to be available on all nodes (striped addresses) or only on a single node (*spotted addresses*). The details of traffic distribution in a cluster depend on the algorithm used, but the same logical entities process the traffic in each case.

Cluster Synchronization

When a node is added to the cluster, the NetScaler configurations and the files (SSL certificates, licenses, DNS, and so on) that are available on the CCO are synchronized on the newly added cluster node. This ensures that the configurations and files are always synchronized on all the nodes of the cluster.

When an existing cluster node rejoins the cluster (after it failed or was intentionally disabled), the cluster checks the configurations available on the node. If there is a mismatch in the configurations available on the rejoined node and on the CCO, the node is synchronized by using one of the following techniques:

- **Full synchronization.** If the difference between configurations exceeds 255 commands, all of the configurations implemented on the CCO are applied to the node that is rejoining the cluster. The node remains operationally unavailable for the duration of the synchronization.
- **Incremental Synchronization.** If the difference between configurations is less than or equal to 255 commands, only the configurations that are not available are applied to the node that is rejoining the cluster. The operational state of the node remains unaffected.

The configurations that are performed on the CCO through the cluster IP address, are automatically propagated to the cluster nodes. Since, the cluster configurations are based on a quorum of the available nodes, a command (that is executed on the cluster

IP address) can be propagated to the other cluster nodes only when a majority of the nodes are in synch. If a majority of the nodes are not in synch or are in the process of synchronizing, they cannot accept new commands and therefore commands are not propagated till the synchronization is complete.

Cluster Connections

To identify the node to which an interface belongs, the standard NetScaler interface-naming convention is prefixed with a node ID. That is, the interface identifier c/u, where c is the controller number and u is the unit number, becomes n/c/u, where n is the node ID.

For example, in the following figure, interface 1/2 of Node 0 is represented as 0/1/2, interface 1/1 of Node 1 is represented as 1/1/1, and interface 1/4 of Node 2 is represented as 2/1/4.

Figure 1-2. Network interfaces naming convention in a cluster

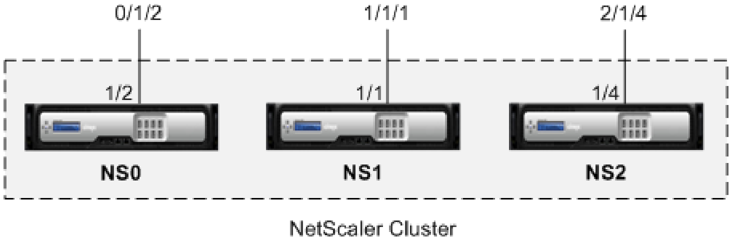

The cluster communicates with the client through the physical connections between the cluster node and the client-side connecting device. The logical grouping of these physical connections is called the *client data plane*. Similarly, the cluster communicates with the server through the physical connections between the cluster nodes and the server-side connecting device. The logical grouping of these physical connections is called the *server data plane*.

In addition to communicating with the client and server through the client data plane and server data plane, respectively, the cluster nodes communicate with each other by using the *cluster backplane*. The backplane, which includes the physical connections from each cluster node and the backplane switch, is the backbone of the cluster system.

Figure 1-3. Cluster communication interfaces

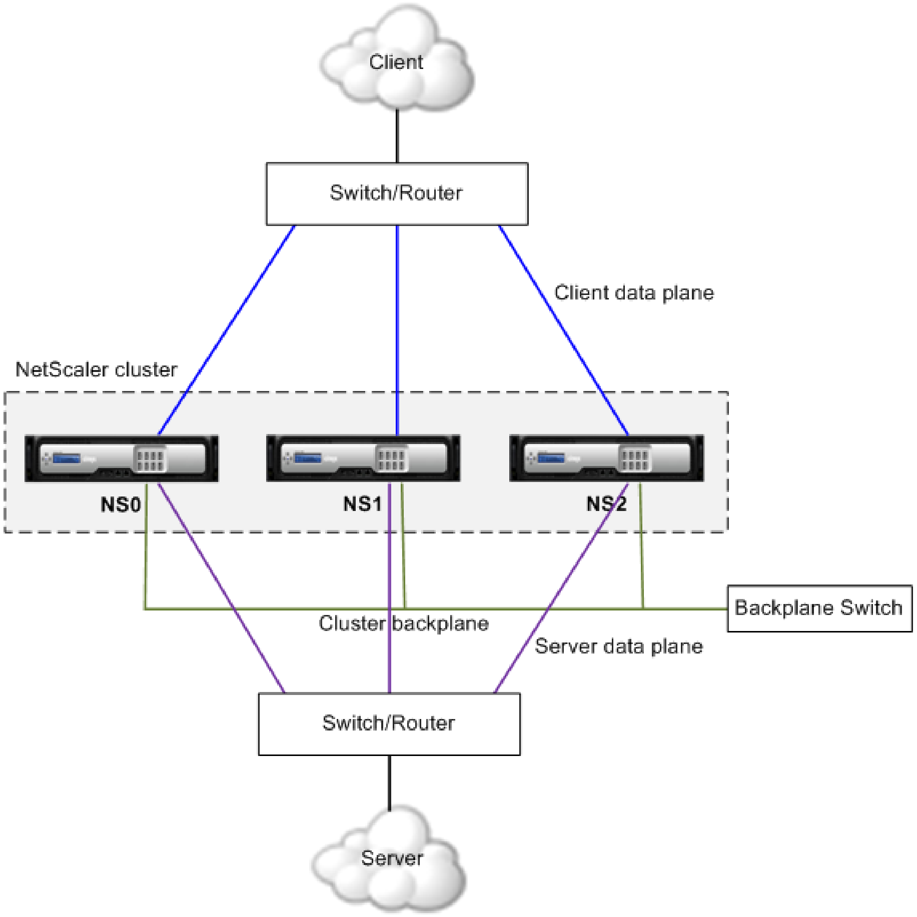

The above figure shows the logical grouping of the physical connections to form the client data plane, server data plane, and cluster backplane.

Striped and Spotted IP Addresses

In a clustered deployment, VIP and SNIP addresses, can be striped or spotted.

- A *striped IP address* is active on all the nodes of the cluster. IP addresses configured on the cluster without specifying an owner node are active on all the cluster nodes.
- A *spotted IP address* is active on and owned exclusively by one node. IP addresses configured on the cluster by specifying an owner node are active only on the node that is specified as the owner.

The following figure shows striped and spotted IP addresses in a three-node cluster.

Figure 1-4. Three-node cluster with striped and spotted IP addresses

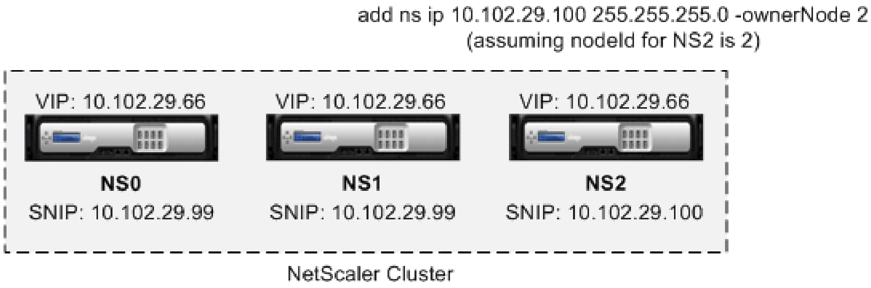

In the above figure, the VIP address 10.102.29.66 is striped on all the cluster nodes, and SNIP address 10.102.29.99 is striped on NS0 and NS1. NS2 has a spotted SNIP address.

The following table shows the NetScaler-owned IP addresses that can be striped or spotted:

Table 1-2. Striped and Spotted IP addresses

NetScaler-owned IP addresses	Striped IP addresses	Spotted IP addresses
NSIP	No	Yes
Cluster IP address	No	No
VIP	Yes	No
SNIP	Yes	Yes (recommended)

Note:

- The cluster IP address is not a striped or spotted IP address. It is a floating IP address that is owned by the CCO, which is not a fixed node.
- Citrix recommends that you use only spotted IP addresses. You can use striped IP addresses only if there is a shortage of IP addresses. The use of striped IP addresses can result in ARP flux issues.

Traffic Distribution

The NetScaler cluster uses Equal Cost Multiple Path (ECMP) or Cluster Link Aggregation Group (CLAG) traffic distribution mechanisms to determine the node that receives the traffic (the *flow receiver*) from the external connecting device. Each of these mechanisms uses a different algorithm to determine the flow receiver. The flow receiver then uses internal cluster logic to determine the node that process the traffic (the *flow processor*).

Note: The flow receiver and flow processor must be nodes capable of serving traffic.

Figure 1-5. Traffic distribution in a cluster

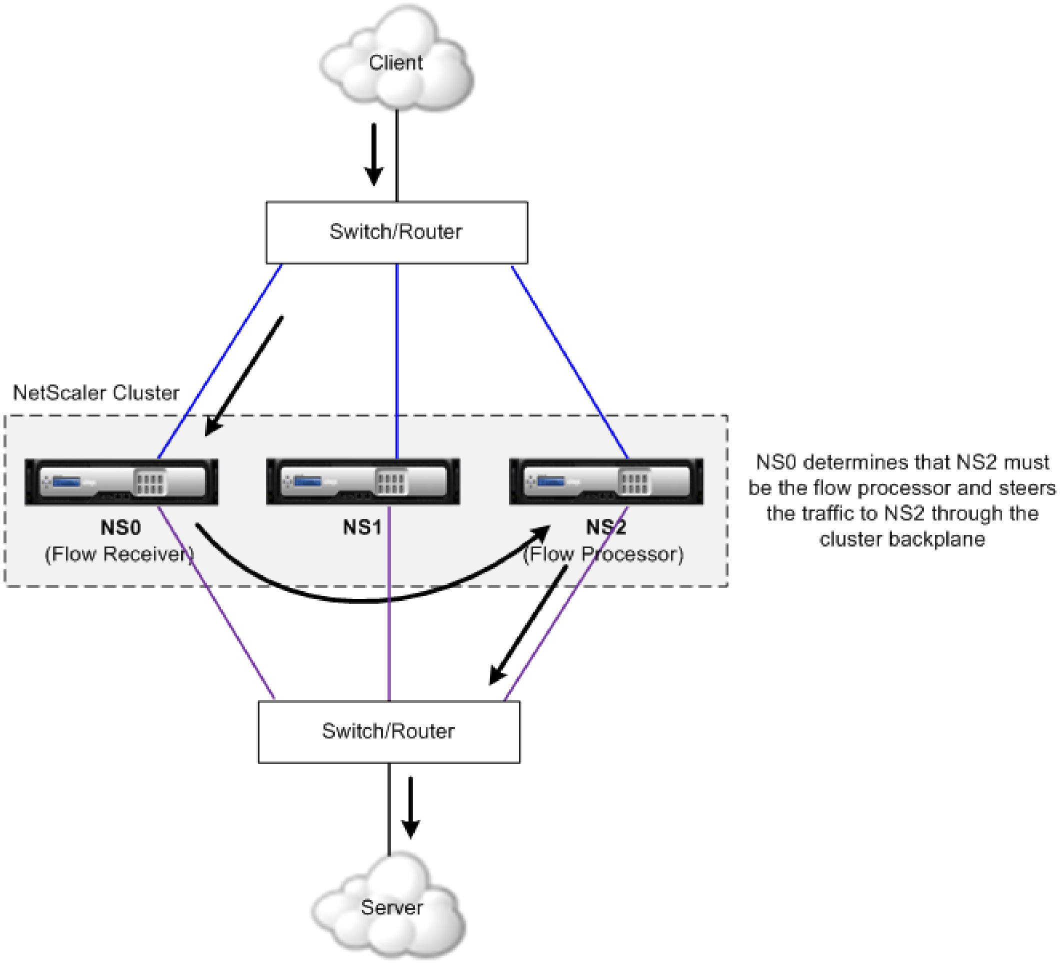
Client
Switch/Router
NetScaler Cluster
NS0
(Flow Receiver)
NS1
NS2
(Flow Processor)
NS0 determines that NS2 must be the flow processor and steers the traffic to NS2 through the cluster backplane
Switch/Router
Server

The above figure shows a client request flowing through the cluster. The client sends a request to a striped virtual IP (VIP) address. A traffic distribution mechanism configured on the client data plane selects one of the cluster nodes as the flow receiver. The flow receiver receives the traffic, determines the node that must process the traffic, and steers the request to that node through the cluster backplane (unless the flow receiver selects itself as the flow processor).

The flow processor establishes a connection with the server. The server processes the request and sends the response to the subnet IP (SNIP) address that sent the request to the server.

- If the SNIP address is a striped IP address, the traffic distribution mechanism configured on the server data plane selects one of the cluster nodes (that owns the SNIP address) as the flow receiver. The flow receiver receives the traffic, determines the flow processor, and steers the request to the flow processor through the cluster backplane.

- If the SNIP address is a spotted IP address, the node that owns the SNIP address receives the response from the server.

In an asymmetric cluster topology (all cluster nodes are not connected to the external switch), you must use *linksets* either exclusively or combined with ECMP or CLAG. For more information, see Using Linksets on page 30.

Cluster and Node States

Cluster-node classification includes three types of states: admin state, operational state, and health.

Admin State. An admin state is configured when you add the node to the cluster. It indicates the purpose of the node, which can be in one of the following states:

- **ACTIVE.** Nodes in this state serve traffic if they are operational and healthy.
- **PASSIVE.** Nodes in this state do not serve traffic but are in sync with the cluster. These nodes are useful during maintenance activity, because they can be upgraded without removing the node from the cluster.
- **SPARE.** Nodes in this state do not serve traffic but are in sync with the cluster. Spare nodes act as backup nodes for the cluster. If one of the ACTIVE nodes becomes unavailable, the operational state of one of the spare nodes becomes ACTIVE, and that node starts serving traffic.

Operational State. When a node is part of a cluster, its operational state can change to ACTIVE, INACTIVE, or UNKNOWN. There are a number of reasons for a node being in INACTIVE or UNKNOWN state. Review the `ns.log` file or error counters to help determine the exact reason.

Health State. Depending on its health, a node can either be UP or NOT UP. To view the reasons for a node being in NOT UP state, run the **show cluster node** command for that node from the cluster IP address.

Only nodes that have the admin state as ACTIVE, operational state as ACTIVE, and health status as UP can serve traffic. A cluster is functional only when a minimum of (n/2 +1) nodes, where n is the number of cluster nodes, are able to serve traffic.

Setting up a NetScaler Cluster

To setup a NetScaler cluster, begin by setting up the cluster backplane. Then, you create the cluster by adding the first node to the cluster, which becomes the initial configuration coordinator (CCO), and by assigning a cluster IP address to that node. Once the cluster IP address is defined on the CCO, you can add more nodes to the cluster.

Every appliance that you want to add to the cluster must:

- Be NetScaler nCore appliances. Clustering of NetScaler Classic appliances is not supported.
- Be of the same platform type (physical appliance or VPX instance).

- Be of the same hardware type (for physical appliances).
- Be on the same subnet.
- Have the cluster license file.
- Have the same licenses (Standard, Enterprise, or Platinum, and any add-on licenses).
- Be of the same software version and build.
- Be initially configured and connected to a common client-side and server-side network.

Only appliances that satisfy all the above criteria can be part of a NetScaler cluster.

Setting up the Cluster Backplane

The nodes in a cluster communicate with each other through the cluster backplane. The backplane is a set of connections in which one interface of each node is connected to a common switch, which is called the *cluster backplane switch*. Each node of the cluster uses a special MAC address to communicate with other nodes through the cluster backplane.

Note: In a cluster of VPX appliances that are deployed on a XenServer (with MAC spoofing enabled), the NIC (XenServer Vswitch) can drop packets sent on the backplane. So, you must make sure that MAC spoofing is disabled on the XenServer.

Note: You must make sure that the cluster backplane switch supports packets larger than 1500 bytes.

Points to remember:

- Do not use the appliance's management interface (0/1) as the backplane interface.
- Interfaces used for the backplane must not be used for the client data plane or server data plane.
- The backplane interfaces of all nodes of a cluster must be connected to the same switch and bound to the same L2 VLAN. The backplane interfaces, by default, have presence on all L3 VLANs configured on the cluster.
- If you have multiple NetScaler clusters with the same cluster instance ID, make sure that the backplane interfaces of each cluster are bound to a different VLAN.
- Citrix recommends that you dedicate a separate switch only for the backplane, so that large amounts of traffic are handled seamlessly.
- The backplane interface is always monitored, regardless of the HA monitoring settings of that interface.

To set up the cluster backplane, do the following for every node

1. Identify the network interface that you want to use for the backplane.

2. Connect an Ethernet or optical cable from the selected network interface to the cluster backplane switch.

For example, to use interface 1/2 as the backplane interface for node 4, connect a cable from the 1/2 interface of node 4 to the backplane switch.

> **Note:** You can configure a link aggregate (LA) channel to optimize the throughput of the cluster backplane.

Creating a NetScaler Cluster

To create a cluster, you must create a cluster instance and configure a cluster IP address on the first appliance that you add to the cluster. This node is called the configuration coordinator (CCO). All cluster configurations are performed on this node, by accessing it through the cluster IP address. The CCO is not fixed to one specific cluster node. It can change over time. For example, if the CCO goes down, the cluster elects one of the other nodes as the new CCO, which then owns the cluster IP address.

When you add the cluster instance, the **clear ns config extended** command is internally executed on that node. In addition, the SNIP addresses and all VLAN configurations (except the default VLAN and NSVLAN) are cleared from the node.

> **Note:** Before you create the cluster, make sure that you have set up the backplane interface for that node.

To create a cluster by using the NetScaler command line

> **Note:** The following commands include only the mandatory parameters. For more information about the CLI commands, see the man pages available for each command. Type `man <command syntax>`. For example, to get the man page for the add cluster instance command, type `man add cluster instance`.

1. Log on to a NetScaler appliance (for example, an appliance with NSIP address 10.102.29.60) that you intend to add to the cluster.
2. Add a cluster instance. The cluster instance is an entity that identifies the cluster.

   add cluster instance <clId>

   where **clId** is a unique number that identifies the cluster. Minimum value: 1. Maximum value: 16.

   > **Note:** Make sure that the cluster instance ID is unique within a LAN.

3. Add the NetScaler appliance to the cluster.

   add cluster node <nodeId> <IPAddress> [-state <state>] [-backplane <interface_name>]

   where,

- **nodeId** is a unique number that identifies the appliance on the cluster. Each node must have a different node ID. Minimum value: 0. Maximum value: 31.
- **IPAddress** is the IP address of the NetScaler appliance. Only IPv4 addresses are supported.
- **state** is the configured state of the cluster node. Possible values: ACTIVE, PASSIVE, SPARE. Default: PASSIVE.

  > **Note:** If you want to perform node-specific configurations, such as adding spotted IP addresses, before the node serves traffic, set the state to PASSIVE (default state). After performing the node-specific configurations, change the node state to ACTIVE by using the **set cluster node** command.

- **backplane** is the backplane interface of the node. For example, if node 0 uses interface 1/1, the value of this parameter is 0/1/1.

Example

```
add cluster node 0 10.102.29.60 -state PASSIVE -backplane 0/1/1
```

4. Add the cluster IP address (for example, 10.102.29.61) on this node.

   add ns ip <IPAddress> <netmask> -type clip

   where,

   - **IPAddress** is the cluster IP address of the NetScaler cluster. Only IPv4 addresses are supported.
   - **netmask** is the subnet mask for the cluster IP address. The value must be 255.255.255.255.

   Example

   ```
   add ns ip 10.102.29.61 255.255.255.255 -type clip
   ```

5. Enable the cluster instance to create the cluster.

   enable cluster instance <clId>

   where **clId** is the number that identifies the cluster instance that must be enabled.

6. Save the configuration.

   save ns config

7. Warm reboot the appliance.

   reboot -warm

Verify the cluster configurations by using the **show cluster instance** command. The output of the command must display the NSIP address of the CCO as a node of the cluster.

To create a cluster by using the configuration utility

1. Log on to a NetScaler appliance (for example, an appliance with NSIP address 10.102.29.60) that you intend to add to the cluster.
2. In the navigation pane, expand **System**, and then click **Cluster**.

3. In the details pane, under **Get Started**, click **Manage Cluster**.
4. In the **Cluster Configuration** dialog box, set the following parameters:
   - **Cluster instance id** - A unique number that identifies the cluster. Minimum value: 1. Maximum value: 16.
   - **Cluster IP address** - The IP address of the NetScaler cluster. Only IPv4 addresses are supported.
   - **Backplane** - The backplane interface of the node. For example, if node 0 uses interface 1/1, the value of this parameter is 1/1.
5. Click **Create**.
6. In the **Configure cluster instance** dialog box, make sure that the **Enable cluster instance** check box is selected.
7. In the **Cluster Nodes** pane, select the node and click **Open**.
8. In the **Configure Cluster Node** dialog box, set the **State**.
9. Click **OK**, and then click **Save**.
10. Warm reboot the appliance.

Adding a Node to the Cluster

You can seamlessly scale the size of a cluster to include a maximum of 32 nodes. When an appliance is added to the cluster, the licenses on that appliance are checked against the licenses available on the CCO. If the licenses match, the appliance is added to the cluster. The existing configurations of the node are cleared, and the cluster configurations are synchronized with the node. There can be an intermittent drop in traffic while the synchronization is in progress.

To add a node to a cluster, you must first configure the node on the cluster (adding the node) and then, configure the cluster on the node (joining the cluster).

If you use the NetScaler command line, first log on to the cluster IP address to add the node. Then, log on to that node and join the node to the cluster. If you use the configuration utility, then you must only log on to the cluster IP address to add the node. The newly added node is automatically joined to the cluster. Alternatively, you can add the node from the command line and use the configuration utility to join the node to the cluster.

Note:

- Before you add the node, make sure that you have set up the backplane interface for that node.
- When you add a new node to a cluster that has only spotted IPs, the synchronization happens before spotted IP addresses are assigned to that node. In such cases, L3 VLAN bindings and static routes can be lost. To avoid this loss, either add a striped IP or add the L3 VLAN bindings and static routes on the NSIP of the newly added node.

- When a NetScaler appliance with pre-configured link aggregate (LA) channel is added to a cluster, the LA channels continue to exist in the cluster environment. The LA channel is renamed from LA/x to nodeId/LA/x, where LA/x is the LA channel identifier.

To add a node to the cluster by using the NetScaler command line

1. Log on to the cluster IP address and do the following:

   a. Add the NetScaler appliance (for example, 10.102.29.70) to the cluster.

   add cluster node <nodeId> <IPAddress> [-state <state>] [-backplane <interface_name>]

   where,

   - **nodeId** is a unique integer that identifies the appliance on the cluster. Each node must have a different node ID. Minimum value: 0. Maximum value: 31.
   - **IPAddress** is the IP address of the NetScaler appliance. Only IPv4 addresses are supported.
   - **state** is the configured state of the cluster node. Possible values: ACTIVE, PASSIVE, SPARE. Default: PASSIVE.

     Note: If you want to perform node-specific configurations, such as adding spotted IP addresses, before the node serves traffic, set the state to PASSIVE (default state). After performing the node-specific configurations, change the node state to ACTIVE by using the **set cluster node** command.

   - **interface_name** is the backplane interface of the node. For example, if node 1 uses interface 1/1, the value of this parameter is 1/1/1.

   Example

   ```
   add cluster node 1 10.102.29.70 -state PASSIVE -backplane
   1/1/1
   ```

   b. Save the configuration.

   save ns config

2. Log on to the newly added node (for example, 10.102.29.70) and do the following:

   a. Join the node to the cluster.

   join cluster -clip <ip_addr> -password <password>

   where,

   - **clip** is the IP address of the NetScaler cluster. Only IPv4 addresses are supported.
   - **password** is the nsroot password of the CCO.

Example

```
join cluster -clip 10.102.29.61 -password nsroot
```

b. Save the configuration.

   save ns config

c. Warm reboot the appliance.

   reboot -warm

To add a node to the cluster by using the configuration utility

1. Log on to the cluster IP address.
2. In the navigation pane, expand **System**, and then click **Cluster**.
3. In the details pane, under **Get Started**, click **Manage Cluster**.
4. Click **Add** to add the new node (for example, 10.102.29.70).
5. In the **Create Cluster Node** dialog box, set the following parameters:
   - **Node Id** - A unique integer that identifies the appliance on the cluster. Each node must have a different node ID. Minimum value: 0. Maximum value: 31.
   - **IP Address** - The IP address of the NetScaler appliance. Only IPv4 addresses are supported.
   - **Backplane** - The backplane interface of the node. For example, if node 1 uses interface 1/1, the value of this parameter is 1/1.
   - **State** - The configured state of the cluster node. Possible values: ACTIVE, PASSIVE, SPARE. Default: PASSIVE.

   > **Note:** If you want to perform node-specific configurations, such as adding spotted IP addresses, before the node serves traffic, set the state to PASSIVE (default state). After performing the node-specific configurations, change the node state to ACTIVE.
6. Click **Create**. A dialog box informs you that the appliance will be warm reboot. Click **Yes** to confirm.

To join a previously added node to the cluster by using the configuration utility

If you have used the NetScaler command line to add a node to the cluster, but have not joined the node to the cluster, you can use the following procedure to join the node to the cluster.

1. Log on to the node that you want to join to the cluster (for example, 10.102.29.70).
2. In the navigation pane, expand **System**, and then click **Cluster**.
3. In the details pane, under **Get Started**, click **Join Cluster**.
4. In the **Join to existing cluster** dialog box, set the following parameters:

- **Cluster IP** - The IP address of the NetScaler cluster. Only IPv4 addresses are supported.
- **Password** - The nsroot password of the CCO.

5. Click **OK**.

Removing a Cluster Node

Removing a node from a cluster is a two-step process:

1. Remove the reference to the cluster instance from the node. This command internally executes the **clear ns config extended** command on that node. In addition, the SNIP addresses and all VLAN configurations (except default VLAN and NSVLAN) are cleared from the node.
2. Remove the node from the cluster.

Note:

- When you remove a node that is the CCO, any current cluster IP address sessions are invalidated. Another cluster node is made the CCO and the cluster IP address is assigned to that node. You must start a new session with the cluster IP address.
- To delete the cluster (and all the nodes), you must remove each node individually. On removing the last node the cluster IP address(es) are deleted.

To remove a cluster node by using the NetScaler command line

1. Log on to the node that you want to remove from the cluster and do the following:
   a. Remove the reference to the cluster instance.

      rm cluster instance <clId>

      where <clId> is the integer that identifies the cluster from which the node is to be removed.

   b. Save the configuration.

      save ns config

   Note: To remove the last node of a cluster, you must only remove the cluster instance from that node. The node is automatically removed from the cluster.

2. Log on to the cluster IP address and do the following:
   a. Remove the node from which you removed the cluster instance.

      rm cluster node <nodeId>

      where <nodeId> is the integer that identifies the node that you are removing.

   b. Save the configuration.

save ns config

Note: Make sure you do not run the **rm cluster node** command from the local node as this results in inconsistent configurations between the CCO and the node.

To remove a cluster node by using the configuration utility

1. Log on to the cluster IP address.
2. In the navigation pane, expand **System**, and then click **Cluster**.
3. In the details pane, under **Get Started**, click **Manage Cluster**.
4. Select the node that you want to remove from the cluster, and then click **Remove**.
5. Click **OK**.

Viewing the Details of a Cluster

You can view the details of the cluster instance and the cluster nodes from the cluster IP address.

To view details of a cluster instance by using the NetScaler command line

At the NetScaler command prompt of the cluster IP address, type:

sh cluster instance <clId>

where <clId> is the integer that identifies the cluster instance whose details you want to view.

```
> show cluster instance 1
1)Cluster ID: 1
  Dead Interval: 3 secs
  Hello Interval: 200 msecs
  Preemption: DISABLED
  Propagation: ENABLED
  Cluster Status: ENABLED(admin), ENABLED(operational), UP

  Member Nodes:
  Node ID    Node IP    Health Admin State Operation State
  -------    -------    ------ ----------- ---------------
1)   0    10.102.29.60*   UP      ACTIVE      ACTIVE(CCO)
2)   1    10.102.29.70    UP      ACTIVE      ACTIVE
Done
```

Note: Executing this command from the NSIP address of a non-CCO node, displays the status of the cluster on this node.

To view details of a cluster node by using the NetScaler command line

At the NetScaler command prompt of the cluster IP address, type:

sh cluster node <nodeId>

where <nodeId> is the integer that identifies the node whose details you want to view.

```
> show cluster node 1
1) Node ID: 1
        IP:                   10.102.29.70
        Backplane:            1/1/1
        Health:               UP
        Admin state:          ACTIVE
        Operational State:    ACTIVE
        Sync State:           ENABLED
Done
```

To view details of a cluster instance by using the configuration utility

1. Log on to the cluster IP address.
2. In the navigation pane, expand **System**, and then click **Cluster.**
3. In the details pane, under **Get Started**, click **Manage Cluster**.
4. In the **Configure Cluster Instance** dialog box, view the details of the cluster.

To view details of a cluster node by using the configuration utility

1. Log on to the cluster IP address.
2. In the navigation pane, expand **System**, click **Cluster**, and then click **Nodes**.
3. In the **Cluster Nodes** list, view the node details. To get a more detailed view of the node, click the node.

Distributing Traffic Across Cluster Nodes

After you have created the NetScaler cluster and performed the required configurations, you must deploy Equal Cost Multiple Path (ECMP) or Cluster Link Aggregation Group (CLAG) on the client data plane (for client traffic) or server data plane (for server traffic). These mechanisms distribute external traffic across the cluster nodes.

Using Equal Cost Multiple Path (ECMP)

With the Equal Cost Multiple Path (ECMP) mechanism, the router has equal-cost routes to VIP addresses with the next hops as the active nodes of the cluster. The router uses a stateless hash-based mechanism to distribute traffic across the routes.

Note: Routes are limited to the maximum number of ECMP routes supported by the upstream router.

To use ECMP, you must first enable the required routing protocol (OSPF, RIP, or BGP) on the cluster IP address. You must bind the interfaces and the spotted IP address (with dynamic routing enabled) to a VLAN. Configure the selected routing protocol and redistribute the kernel routes on the ZebOS by using the vtysh shell.

You must perform similar configurations on the cluster IP address and on the external connecting device.

You must have detailed knowledge of routing protocols to use ECMP. For more information, see the "Configuring Dynamic Routes" chapter of the *Citrix NetScaler Networking Guide*. For a link to the guide, see the Documentation Library on page 65.

> **Note:** Make sure the licenses on the cluster support dynamic routing, otherwise ECMP traffic distribution does not work. The standard NetScaler license, for example, does not support dynamic routing.

Figure 1-6. ECMP topology

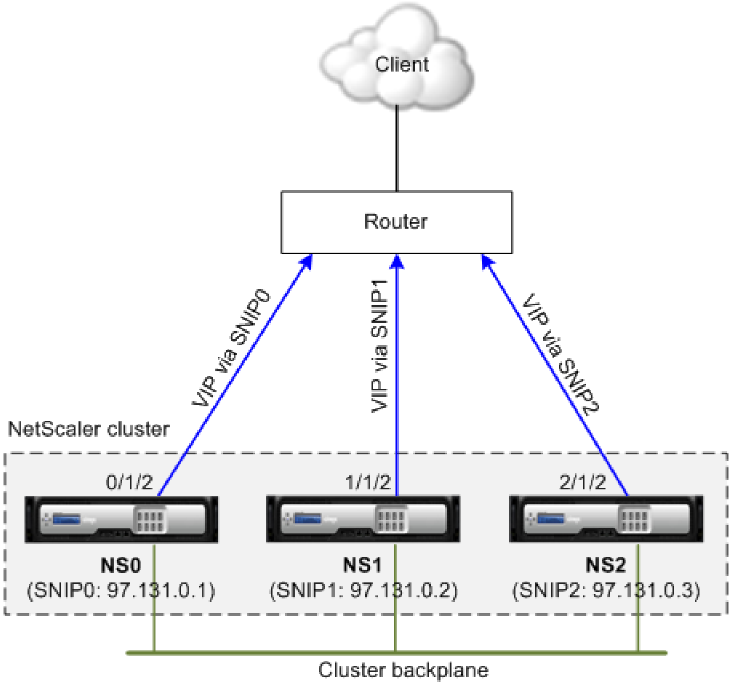

As seen in the above figure, the ECMP router can reach the VIP address via SNIP0, SNIP1, or SNIP2.

To configure ECMP on the NetScaler cluster by using the NetScaler command line

1. Log on to the cluster IP address.
2. Enable the routing protocol (OSPF, RIP, or BGP).

   enable ns feature <routing protocol>

   Example: To enable the OSPF routing protocol.

   ```
   enable ns feature ospf
   ```

3. Add a VLAN.

   add vlan <vlan id>

   Example

   ```
   add vlan 97
   ```

4. Bind the interfaces of the cluster nodes to the VLAN.

   bind vlan <vlan id> -ifnum <interface_name>

   Example

   ```
   bind vlan 97 -ifnum 0/1/2 1/1/2 2/1/2
   ```

5. Add a spotted SNIP address on each node and enable dynamic routing on it.

   add ns ip <SNIP> <netmask> -ownerNode <node id> -dynamicRouting ENABLED

   Example

   ```
   add ns ip 97.131.0.1 255.0.0.0 -ownerNode 0 -dynamicRouting
   ENABLED -type SNIP
   add ns ip 97.131.0.2 255.0.0.0 -ownerNode 1 -dynamicRouting
   ENABLED -type SNIP
   add ns ip 97.131.0.3 255.0.0.0 -ownerNode 2 -dynamicRouting
   ENABLED -type SNIP
   ```

6. Bind one of the spotted SNIP addresses to the VLAN. When you bind one spotted SNIP address to a VLAN, all other spotted SNIP addresses defined on the cluster in that subnet are automatically bound to the VLAN.

   bind vlan <vlan id> -ipAddress <SNIP> <netmask>

   Example

   ```
   bind vlan 97 -ipAddress 97.131.0.1 255.0.0.0
   ```

   Note: You can use NSIP addresses of the cluster nodes instead of adding SNIP addresses. If so, you do not have to perform steps 3 - 6.

7. Configure the routing protocol on ZebOS using vtysh shell.

   Example: To configure OSPF routing protocol on node IDs 0, 1, and 2.

   ```
   !
   interface vlan97
   !
   router ospf
    owner-node 0
     ospf router-id 97.131.0.1
    exit-owner-node
    owner-node 1
     ospf router-id 97.131.0.2
    exit-owner-node
    owner-node 2
     ospf router-id 97.131.0.3
    exit-owner-node
    redistribute kernel
    network 97.0.0.0/8 area 0
   !
   ```

Using Cluster Link Aggregation Group (CLAG)

A cluster link aggregation group (CLAG), as the name suggests, is a group of interfaces of cluster nodes. It is an extension of NetScaler link aggregation. The only difference is that while link aggregation requires the interfaces to be from the same device, in CLAG, the interfaces are from different nodes of the cluster.

For more information about link aggregation, see the "Configuring Link Aggregation" chapter of the *Citrix NetScaler Networking Guide*. For a link to the guide, see the Documentation Library on page 65.

CLAG can be either static or dynamic.

For example, consider a three-node cluster where all three nodes are connected to the upstream switch. A CLAG channel (CLA/1) is formed by binding interfaces 0/1/2, 1/1/2, and 2/1/2.

Figure 1-7. Cluster Link Aggregation Group topology

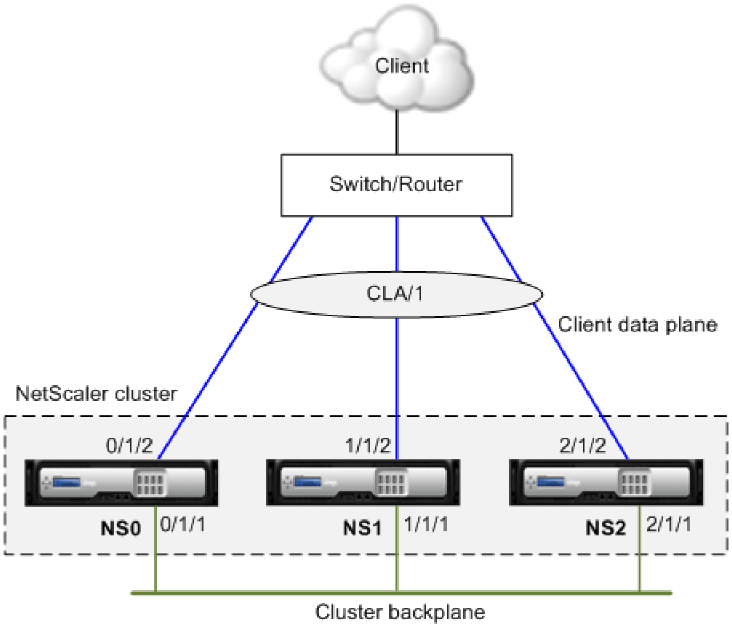

A CLAG channel has the following attributes:

- Each channel has a unique MAC agreed upon by cluster nodes.
- The channel can bind both local and remote nodes' interfaces.
- A maximum of four CLAG channels are supported in a cluster.
- Backplane interfaces cannot be part of a CLAG channel.
- When an interface is bound to a CLAG channel, the channel parameters have precedence over the network interface parameters. A network interface can be bound to one channel only.

Figure 1-8. Traffic distribution flow using CLAG

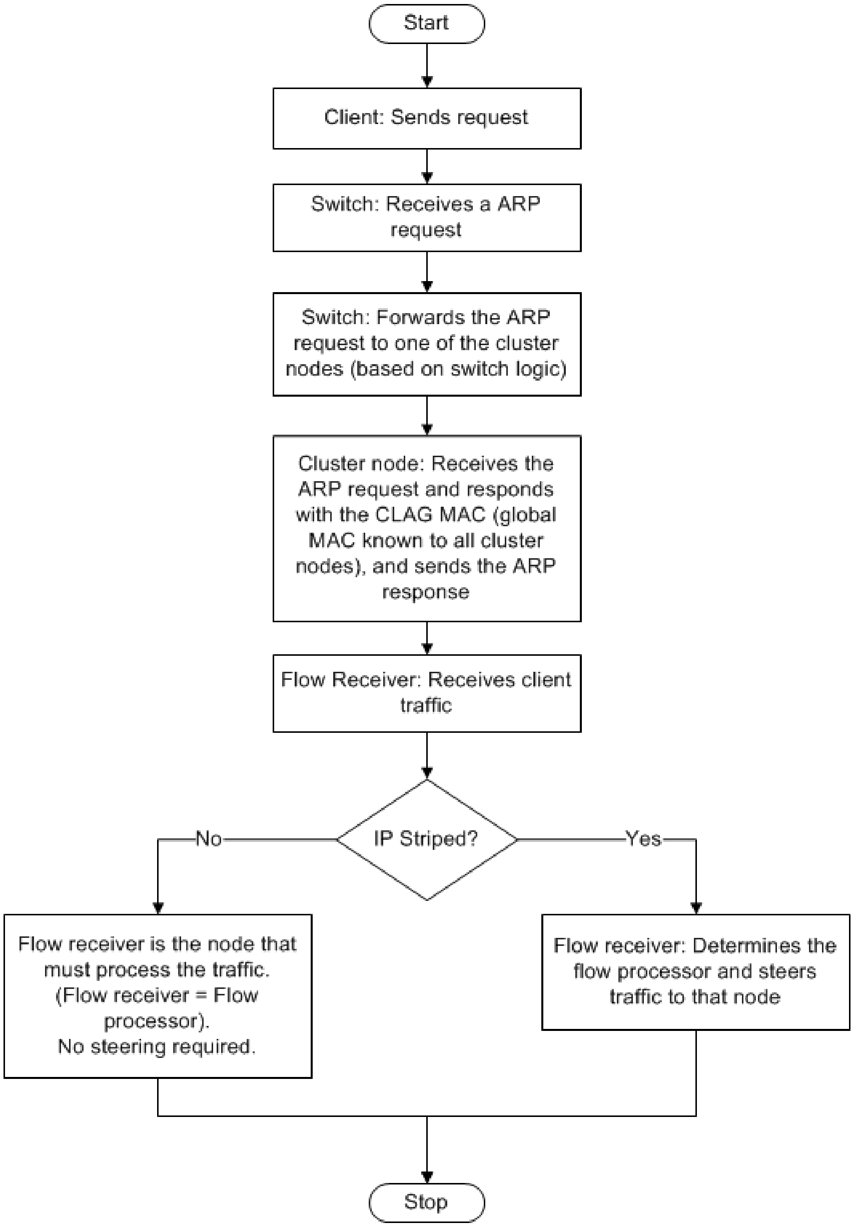
Start
Client: Sends request
Switch: Receives a ARP request
Switch: Forwards the ARP request to one of the cluster nodes (based on switch logic)
Cluster node: Receives the ARP request and responds with the CLAG MAC (global MAC known to all cluster nodes), and sends the ARP response
Flow Receiver: Receives client traffic
IP Striped?
No
Yes
Flow receiver is the node that must process the traffic. (Flow receiver = Flow processor). No steering required.
Flow receiver: Determines the flow processor and steers traffic to that node
Stop

Static Cluster Link Aggregation Group

You must configure a static CLAG channel on the cluster IP address and on the external connecting device. If possible, configure the upstream switch to distribute traffic based on IP address or port instead of MAC address.

For more information about configuring a static LA channel, see the "Configuring Link Aggregation Manually" chapter of the *Citrix NetScaler Networking Guide*. For a link to the guide, see the Documentation Library on page 65.

To configure a static CLAG channel by using the NetScaler command line

1. Log on to the cluster IP address.

   > **Note:** Make sure that you configure the CLAG channel on the cluster IP address before configuring CLAG on the external switch. Otherwise, the switch will forward traffic to the cluster even though the CLAG channel is not configured. This can lead to loss of traffic.

2. Create a CLAG channel.

   add channel <clag channel id> -speed <speed>

   where,

   - <clag channel id> is a unique number that identifies the CLAG channel. Must be of the form CLA/x where x can range from 1 to 4.
   - <speed> is the speed for the member interfaces of the CLAG.

   Example

   ```
   add channel CLA/1 -speed 1000
   ```

   > **Note:** You must not specify the speed as AUTO. Rather, you must explicitly specify the speed as 10, 100, 1000, or 10000. Only interfaces that have the speed matching the <speed> attribute in CLAG are added to the active distribution list.

3. Bind the required interfaces to the CLAG channel. Make sure that the interfaces are not used for the cluster backplane.

   bind channel <clag channel id> <interface_name...>

   where,

   - <clag channel id> identifies the CLAG channel to which you want to bind the interfaces.
   - <interface_name> specifies the interfaces to be bound to the CLAG channel.

   Example

   ```
   bind channel CLA/1 1/1/2 2/1/2 3/1/2
   ```

4. Verify the CLAG channel configurations.

   show channel <clag channel id>

   Example

   ```
   show channel CLA/1
   ```

   > **Note:** You can bind the CLAG channel to a VLAN by using the **bind vlan** command. The interfaces of the CLAG channel are automatically bound to the VLAN.

Dynamic Cluster Link Aggregation Group

Dynamic CLAG uses Link Aggregation Control Protocol (LACP). For more information about configuring a dynamic LA channel, see the "Configuring Link Aggregation by Using the Link Aggregation Control Protocol" chapter of the *Citrix NetScaler Networking Guide*. For a link to the guide, see the Documentation Library on page 65.

You must perform similar configurations on the cluster IP address and on the external connecting device. If possible, configure the upstream switch to distribute traffic based on IP address or port instead of MAC address.

Points to remember:

- Enable LACP (by specifying the LACP mode as either ACTIVE or PASSIVE).

  > **Note:** Make sure the LACP mode is not set as PASSIVE on both the NetScaler cluster and the external connecting device.

- Specify the same LACP key on each interface that you want to be the part of the channel. For creating a CLAG channel, the LACP key can have a value from 5 through 8.

  For example, if you set the LACP Key on interfaces 1/1/2 and 2/1/2 to 5, CLA/1 is created. The interfaces 1/1/2 and 2/1/2 are automatically bound to CLA/1. Similarly, if you set the LACP key to 6, CLA/2 channel is created.

- Specify the LAG type as Cluster.

To configure a dynamic CLAG channel by using the NetScaler command line

On the cluster IP address, for each interface that you want to add to the CLAG channel, type:

set interface <interface id> -lacpMode <lacpMode> -lacpKey <lacpKey> -lagType Cluster

Example: To configure a CLAG channel for 3 interfaces.

```
set interface 0/1/2 -lacpMode active -lacpKey 5 -lagType Cluster
set interface 1/1/2 -lacpMode active -lacpKey 5 -lagType Cluster
set interface 2/1/2 -lacpMode active -lacpKey 5 -lagType Cluster
```

Using Linksets

Linksets must be used when some cluster nodes are not physically connected to the external network. In such a cluster topology, the unconnected cluster nodes use the interfaces specified in the linkset to communicate with the external network through the cluster backplane. Linksets are typically used in scenarios when the connecting devices have insufficient ports to connect the cluster nodes.

Linksets must be configured only on the cluster IP address.

For example, consider a three-node cluster where the upstream switch has only two ports available. Using linksets, you can connect two nodes to the switch and leave the third node unconnected. In the following figure, a linkset (LS/1) is formed by binding the interfaces 0/1/2 and 1/1/2. NS2 is the unconnected node of the cluster.

Note: Use linksets for better performance of topologies that require MAC-based Forwarding (MBF).

Figure 1-9. Linksets topology

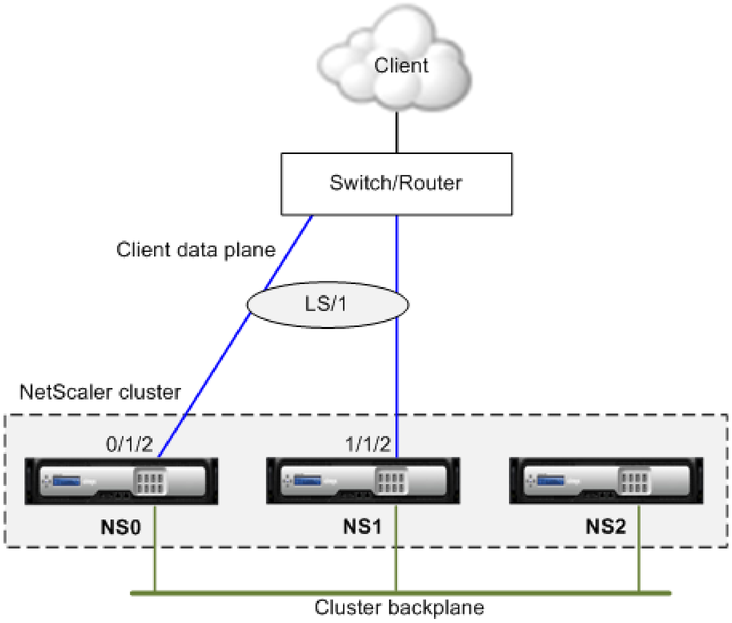

The linkset informs NS2 that it can use interfaces 0/1/2 and 1/1/2 to communicate with the network devices. All traffic to and from NS2 is now routed through interfaces 0/1/2 or 1/1/2.

Figure 1-10. Traffic distribution flow using linksets

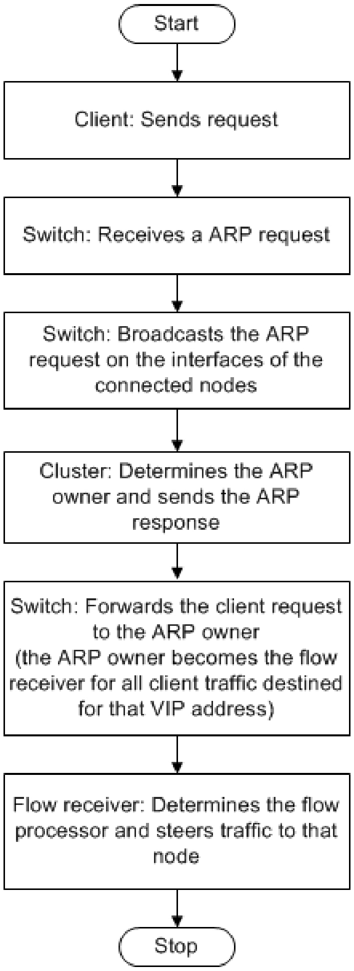
Start
Client: Sends request
Switch: Receives a ARP request
Switch: Broadcasts the ARP request on the interfaces of the connected nodes
Cluster: Determines the ARP owner and sends the ARP response
Switch: Forwards the client request to the ARP owner (the ARP owner becomes the flow receiver for all client traffic destined for that VIP address)
Flow receiver: Determines the flow processor and steers traffic to that node
Stop

To configure a linkset by using the NetScaler command line

1. Log on to the cluster IP address.
2. Create a linkset.

   add linkset <linkset id>

   where <linkset id> is a unique identifier of the linkset. It must be of the form LS/x.

   Example

   ```
   add linkset LS/1
   ```
3. Bind the required interfaces to the linkset. Make sure the interfaces are not used for the cluster backplane.

   bind linkset <linkset id> -ifnum <interface_name...>

   where <interface_name> specifies the interfaces to be bound to the linkset.

Example

```
bind linkset LS/1 -ifnum 0/1/2 1/1/2
```

4. Verify the linkset configurations.

   show linkset <linkset id>

   where <linkset id> is a identifier of the linkset you want to verify.

   Example

   ```
   show linkset LS/1
   ```

> **Note:** You can bind the linkset to a VLAN by using the **bind vlan** command. The interfaces of the linkset are automatically bound to the VLAN.

To configure a linkset by using the configuration utility

1. Log on to the cluster IP address.
2. In the navigation pane, expand **Network**, and then click **Linkset**.
3. In the details pane, click **Add**.
4. In the **Create Linkset** dialog box:
   a. Specify the name of the linkset by setting the **Linkset** parameter.
   b. Specify the **Interfaces** to be added to the linkset and click **Add**. Repeat this step for each interface you want to add to the linkset.
5. Click **Create**, and then click **Close**.

Managing the NetScaler Cluster

After you have created a cluster and configured the required traffic distribution mechanism, the cluster is able to serve traffic. During the lifetime of the cluster, you can perform cluster management tasks such as disabling nodes of a cluster, discovering NetScaler appliances, viewing statistics, synchronizing cluster configurations, cluster files, and the time across the nodes, and upgrading or downgrading the software of cluster nodes.

Disabling a Cluster Node

You can temporarily remove a node from a cluster by disabling the cluster instance on that node. A disabled node is not synchronized with the cluster configurations and is unable to serve traffic.

To disable a cluster node by using the NetScaler command line

At the NetScaler command prompt of the node that you want to disable, type:

disable cluster instance <clId>

where, <clId> identifies the cluster instance that you want to disable.

> **Note:** To disable the cluster, run the **disable cluster instance** command on the cluster IP address.

To disable a cluster node by using the configuration utility

1. Log on to the node that you want to disable.
2. In the navigation pane, expand **System**, and then click **Cluster**.
3. In the details pane, under **Get Started**, click **Manage Cluster**.
4. In the **Configure cluster instance** dialog box, unselect the **Enable cluster instance** check box.
5. Click **OK**.

   > **Note:** To disable the cluster instance on all the nodes, log on to the cluster and perform the above procedure.

Discovering NetScaler Appliances

You can discover NetScaler appliances present in the same subnet as the NSIP address of the CCO. The discovered appliances can then be added to the cluster.

> **Note:** This operation is only available through the configuration utility.

To discover appliances by using the NetScaler configuration utility

1. Log on to the cluster IP address.
2. In the navigation pane, expand **System**, click **Cluster**, and then click **Nodes**.
3. In the details pane, at the bottom of the page, click **Discover NetScalers**.
4. In the **Discover NetScalers** dialog box, set the following parameters:
   - **IP address range** - Specify the range of IP addresses within which you want to discover NetScaler appliances. For example, you can search for all NSIP addresses between 10.102.29.4 to 10.102.29.15 by specifying this option as 10.102.29.4 - 15.
   - **Backplane interface** - Specify the interfaces to be used as the backplane interface. This is an optional parameter. If you do not specify this parameter, you must update it after the node is added to the cluster.
5. Click **OK**.
6. Select the NetScaler appliances that you want to add to the cluster.
7. Click **OK**.

Viewing the Statistics of a Cluster

You can view the statistics of a cluster instance and cluster nodes to evaluate the performance or to troubleshoot the operation of the cluster.

To view the statistics of a cluster instance by using the NetScaler command line

At the NetScaler command prompt of the cluster IP address, type:

stat cluster instance <clId>

```
> stat cluster instance

Cluster Instance Summary
Cluster Size                                       3
Cluster Status                               ENABLED
Cluster Config Coordinator (CCO)        10.102.29.80
Current DFD Sessions                               0
Total Steered Packets                              0
Done
```

To display the statistics of the cluster instance with the error statistics, at the NetScaler command prompt of the cluster IP address, type:

stat cluster instance -detail <clId>

```
> stat cluster instance -detail
Cluster Statistics

Summary
Cluster Size                                       3
Cluster Status                               ENABLED
Cluster Config Coordinator (CCO)        10.102.29.80
Current DFD Sessions                               0
Total Steered Packets                              0

Error Statistics
DFD Dropped Packets                                0
Propagation timeout                                0

Done
```

To view the statistics of a cluster node by using the NetScaler command line

At the NetScaler command prompt of the cluster IP address, type:

stat cluster node

```
> stat cluster node

Cluster Node Summary
NodeID NodeIP       State    Health  Sync State HB Tx HB Rx
0      10.102.29.70 ACTIVE   UP      ENABLED    4489  2247
1      10.102.29.80 ACTIVE   UP      ENABLED    2659  4805
```

```
2      10.102.29.60 INACTIVE UNKNOWN UNKNOWN    7145  0
 Done
```

To display the statistics of an individual cluster node, at the NetScaler command prompt of the cluster IP address, type:

stat cluster node <nodeid>

```
> stat cluster node 1

Node ID : 1
Node IP                                 10.102.29.80
Master State                                  ACTIVE
Health                                            UP
Sync State                                   ENABLED
Heartbeats Sent                                 3025
Heartbeats received                             5537

NNM Statistics
NNM current connections                            7
NNM total transmitted messages                    15
NNM total received messages                       18

Error Statistics
NNM Multicast/Broadcast req err                    0
 Done
```

To view the statistics of a cluster instance by using the configuration utility

1. Log on to the cluster IP address.
2. In the navigation pane, expand **System**, and then click **Cluster**.
3. In the details pane, in the center of the page, click **Statistics**.

To view the statistics of a cluster node by using the configuration utility

1. Log on to the cluster IP address.
2. In the navigation pane, expand **System**, click **Cluster**, and then click **Nodes**.
3. In the details pane, select a node and click **Statistics** to view the statistics of the node. To view the statistics of all the nodes, click **Statistics** without selecting a specific node.

Synchronizing Cluster Configurations

NetScaler configurations that are available on the CCO are synchronized to the other nodes of the cluster when:

- A node joins the cluster
- A node rejoins the cluster

- A new command is executed on the CCO.

Additionally, you can forcefully synchronize the configurations that are available on the CCO (full synchronization) to a specific cluster node. Make sure you synchronize one cluster node at a time, otherwise the cluster can get affected.

To synchronize cluster configurations by using the NetScaler command line

At the NetScaler command prompt of the appliance on which you want to synchronize the CCO configurations, type:

```
force cluster sync
```

For detailed information on the command, see the man page of the command.

To synchronize cluster configurations by using the configuration utility

1. Log on to the appliance on which you want to synchronize the CCO configurations.
2. In the navigation pane, expand **System**, and then click **Cluster.**
3. In the details pane, under **Utilities**, click **Force cluster sync**.
4. Click **OK.**

Synchronizing Cluster Files

The files available on the CCO are called cluster files. These files are automatically synchronized on the other cluster nodes when the node is added to the cluster and periodically, during the lifetime of the cluster. Additionally, you can manually synchronize the cluster files.

The directories and files from the CCO that are synchronized are:

- /nsconfig/ssl/
- /var/netscaler/ssl/
- /var/vpn/bookmark/
- /nsconfig/dns/
- /nsconfig/htmlinjection/
- /netscaler/htmlinjection/ens/
- /nsconfig/monitors/
- /nsconfig/nstemplates/
- /nsconfig/ssh/
- /nsconfig/rc.netscaler
- /nsconfig/resolv.conf

- /nsconfig/inetd.conf
- /nsconfig/syslog.conf
- /nsconfig/snmpd.conf
- /nsconfig/ntp.conf
- /nsconfig/httpd.conf
- /nsconfig/sshd_config
- /nsconfig/hosts
- /nsconfig/enckey
- /var/nslw.bin/etc/krb5.conf
- /var/nslw.bin/etc/krb5.keytab
- /var/lib/likewise/db/
- /var/download/
- /var/wi/tomcat/webapps/
- /var/wi/tomcat/conf/Catalina/localhost/
- /var/wi/java_home/lib/security/cacerts
- /var/wi/java_home/jre/lib/security/cacerts
- /nsconfig/license/
- /nsconfig/rc.conf

To synchronize cluster files by using the NetScaler command line

At the NetScaler command prompt of the cluster IP address, type:

sync cluster files <mode>

where, <mode> specifies the directories or files to be synchronized. Possible values are: all, bookmarks, ssl, htmlinjection, imports, misc, dns, all_plus_misc. Default value: all.

For detailed information on the command, see the man page of the command.

To synchronize cluster files by using the configuration utility

1. Log on to the cluster.
2. In the navigation pane, expand **System**, and then click **Cluster.**
3. In the details pane, under **Utilities**, click **Synchronize cluster files**.
4. In the **Synchronize cluster files** dialog box, select the files to be synchronized in the **Mode** drop-down box.
5. Click **OK**.

Synchronizing Time on Cluster Nodes

The NetScaler cluster uses Precision Time Protocol (PTP) to synchronize the time across cluster nodes. PTP uses multicast packets to synchronize the time. If there are some issues in time synchronization, you must disable PTP and configure Network Time Protocol (NTP) on the cluster.

To enable/disable PTP by using the NetScaler command line

At the NetScaler command prompt of the cluster IP address, type:

```
set ptp -state disable
```

To enable/disable PTP by using the configuration utility

1. Log on to the cluster IP address.
2. In the navigation pane, expand **System**, and then click **Cluster**.
3. In the details pane, under **Utilities**, click **Configure PTP Settings**.
4. In the **Enable/Disable PTP** dialog box, select whether you want to enable or disable PTP.
5. Click **OK**.

Upgrading or Downgrading Software of the Cluster

All cluster nodes must be running the same software version. To upgrade or downgrade the software of a cluster, you must upgrade or downgrade the software on each node, one node at a time.

When the software on a node is upgraded or downgraded, the node is not removed from the cluster. The node continues to be a part of the cluster and serves client traffic uninterrupted, except for the down-time when the node reboots after it is upgraded or downgraded. However, due to software version mismatch among the cluster nodes, configuration propagation is disabled and is enabled only after all the cluster nodes are of the same version.

Since configuration propagation is disabled during upgrading on downgrading a cluster, you cannot perform any configurations through the cluster IP address during this time. However, you can perform node-level configurations through the NSIP address of individual nodes, but you must make sure that you perform the same configurations on all the nodes to maintain them in synch.

> **Note:** You cannot add cluster nodes while upgrading or downgrading the cluster software version.

To upgrade or downgrade the software of the cluster nodes

1. Make sure the cluster is stable and the configurations are synchronized on all the nodes.
2. Upgrade or downgrade the software of the cluster.
   a. Upgrade or downgrade the software of a cluster node. For detailed information about upgrading and downgrading the software of an appliance, see the "Upgrading or Downgrading the System Software" chapter of the *Citrix NetScaler Migration Guide*. For a link to the guide, see the Documentation Library on page 65.

   b. Reboot the appliance.

   c. Repeat the above two steps for each of the other cluster nodes.

   > **Note:** Citrix recommends that you wait for the previous node to become active before upgrading the next node.

Use Cases

This topic provides some use cases for deploying a NetScaler cluster.

- Creating a Two-Node Cluster on page 40
- Migrating an HA Setup to a Cluster Setup on page 41
- Using Cache Redirection in a Cluster on page 42
- Using CLAG with Linksets on page 42
- Backplane on LA Channel on page 43
- Common Interfaces for Client and Server and Dedicated Interfaces for Backplane on page 44
- Common Switch for Client, Server, and Backplane on page 46
- Common Switch for Client and Server and Dedicated Switch for Backplane on page 48
- Multiple Switches for Each Node on page 51
- Different Switch for Every Node on page 52
- Sample Cluster Configurations on page 53

Creating a Two-Node Cluster

A two-node cluster is an exception to the rule that a cluster is functional only when a minimum of (n/2 +1) nodes, where n is the number of cluster nodes, are able to serve traffic. If that formula were applied to a two-node cluster, the cluster would fail if one node went down (n/2 +1=2).

A two-node cluster is functional even if only one node is able to serve traffic.

Creating a two node cluster is the same as creating any other cluster. You must add one node as the configuration coordinator and the other node as the other cluster node.

> **Note:** Incremental configuration synchronization is not supported in a two-node cluster. Only full synchronization is supported.

Migrating an HA Setup to a Cluster Setup

An existing high availability (HA) setup can be migrated to a cluster setup by removing the appliances from the HA setup and then creating the NetScaler cluster. For example, consider an HA setup with NSIP addresses 10.102.97.131 and 10.102.97.132.

To convert a HA setup to cluster setup by using the NetScaler command line

1. Log on to each HA node and remove it from the HA setup.

   rm HA node <nodeId>

   Example

   ```
   rm HA node 1
   ```

2. Go to the shell on one of the HA nodes and copy the `ns.conf` to another .conf file (for example, `ns_backup.conf`).

3. Edit the new configuration file as follows:

   - Remove all features that are not supported by a cluster. For the list of unsupported features, see NetScaler Features Supported by a Cluster on page 7.
   - Remove the configurations that have interfaces, or update the interface names from the c/u convention to the n/c/u convention.

4. On both the nodes, identify the network interfaces to be used for the cluster backplane.

5. Set up one of the nodes (for example, 10.102.97.131) as the CCO node. For detailed instructions, see Creating a NetScaler Cluster on page 17.

6. Log on to the cluster IP address and apply configurations from the backup configuration file.

   batch -f <fileName>

   Example

   ```
   batch -f ns_backup.conf
   ```

7. Save the configuration.

   save ns config

8. Add the other node to the cluster. For detailed instructions, see Adding a Node to the Cluster on page 19.

The appliances of the HA setup are migrated to a cluster setup.

Using Cache Redirection in a Cluster

Cache redirection in a NetScaler cluster works in the same way as it does on a standalone NetScaler appliance. The only difference is that the configurations are done on the cluster IP address. For more information on Cache Redirection, see the "Cache Redirection" chapter of the *Citrix NetScaler Traffic Management Guide*. For a link to the guide, see the Documentation Library on page 65.

Points to remember when using cache redirection in transparent mode:

- Before configuring cache redirection, make sure that you have connected all nodes to the external switch and that you have linksets configured. Otherwise, client requests will be dropped.
- When MAC mode is enabled on a load balancing virtual server, make sure the MBF mode is enabled on the cluster (using **enable ns mode MBF** command). Otherwise, the requests are sent to origin server directly instead of being sent to the cache server.

Using CLAG with Linksets

In an asymmetric cluster topology, some cluster nodes are not connected to the upstream network. In such a case, you must use linksets. To optimize the performance, you can bind the interfaces that are connected to the switch as a CLA channel and then bind the CLA channel to a linkset.

To understand how a combination of CLAG and linksets can be used, consider a three-node cluster for which the upstream switch has only two ports available. You can connect two of the cluster nodes to the switch and leave the other node unconnected.

Note: Similarly, you can also use a combination of ECMP and linksets in an asymmetric topology.

Figure 1-11. Linksets and Cluster Link Aggregation Group topology

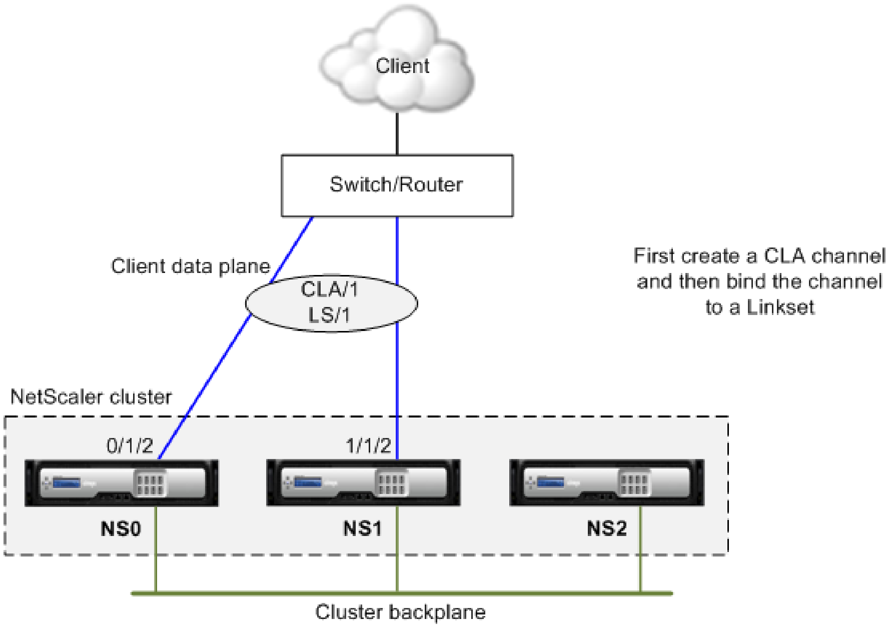

To use CLAG and linksets by using the NetScaler command line

1. Log on to the cluster IP address.
2. Bind the connected interfaces to a CLA channel.

```
add channel CLA/1 -ifnum 0/1/2 1/1/2
```

3. Bind the CLA channel to a linkset.

```
add linkset LS/1 -ifnum CLA/1
```

Backplane on LA Channel

In this deployment, LA channels are used for the cluster backplane.

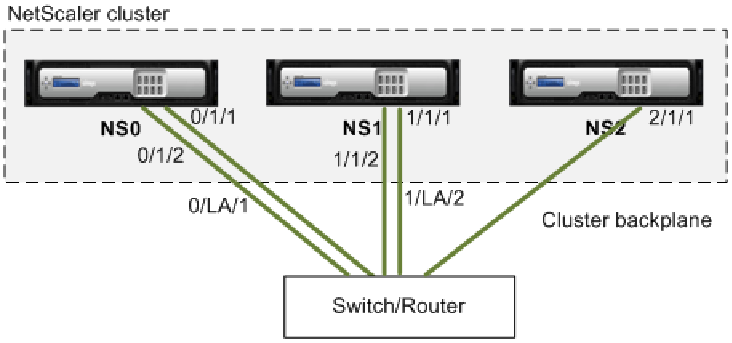

NS0 - nodeId: 0, NSIP: 10.102.29.60

NS1 - nodeId: 1, NSIP: 10.102.29.70

NS2 - nodeId: 2, NSIP: 10.102.29.80

To deploy a cluster with the backplane interfaces as LA channels

1. Create a cluster of nodes NS0, NS1, and NS2.

   a. Log on to the first node that you want to add to the cluster and do the following:

   ```
   create cluster instance 1
   add cluster node 0 10.102.29.60 -state ACTIVE
   enable cluster instance 1
   add ns ip 10.102.29.61 255.255.255.255 -type CLIP
   save ns config
   reboot -warm
   ```

   b. Log on to the cluster IP address and do the following:

   ```
   add cluster node 1 10.102.29.70 -state ACTIVE
   add cluster node 2 10.102.29.80 -state ACTIVE
   ```

   c. Log on to the nodes 10.102.29.70 and 10.102.29.80 to join the nodes to the cluster.

   ```
   join cluster -clip 10.102.29.61 -password nsroot
   save ns config
   reboot -warm
   ```

   As seen in the above commands the interfaces 0/1/1, 1/1/1, and 2/1/1 are configured as the backplane interfaces of the three cluster nodes.

2. Log on to the cluster IP address and do the following:

   a. Create the LA channels for nodes NS0 and NS1.

   ```
   add channel 0/LA/1 -ifnum 0/1/1 0/1/2
   add channel 1/LA/2 -ifnum 1/1/1 1/1/2
   ```

   b. Configure the backplane for the cluster nodes.

   ```
   set cluster node 0 -backplane 0/LA/1
   set cluster node 1 -backplane 1/LA/2
   set cluster node 2 -backplane 2/1/1
   ```

Common Interfaces for Client and Server and Dedicated Interfaces for Backplane

This is a one-arm deployment of the NetScaler cluster. In this deployment, the client and server networks use the same interfaces to communicate with the cluster. The cluster backplane uses dedicated interfaces for inter-node communication.

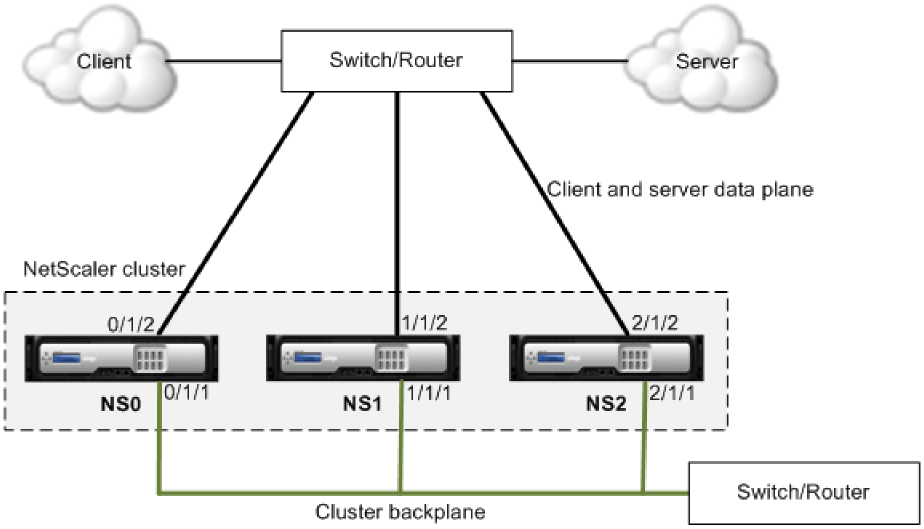

NS0 - nodeId: 0, NSIP: 10.102.29.60

NS1 - nodeId: 1, NSIP: 10.102.29.70

NS2 - nodeId: 2, NSIP: 10.102.29.80

To deploy a cluster with a common interface for the client and server and a different interface for the cluster backplane

1. Create a cluster of nodes NS0, NS1, and NS2.

   a. Log on to the first node that you want to add to the cluster and do the following:

   ```
   create cluster instance 1
   add cluster node 0 10.102.29.60 -state ACTIVE -backplane
   0/1/1
   enable cluster instance 1
   add ns ip 10.102.29.61 255.255.255.255 -type CLIP
   save ns config
   reboot -warm
   ```

   b. Log on to the cluster IP address and do the following:

   ```
   add cluster node 1 10.102.29.70 -state ACTIVE -backplane
   1/1/1
   add cluster node 2 10.102.29.80 -state ACTIVE -backplane
   2/1/1
   ```

   c. Log on to the nodes 10.102.29.70 and 10.102.29.80 to join the nodes to the cluster.

   ```
   join cluster -clip 10.102.29.61 -password nsroot
   save ns config
   reboot -warm
   ```

   As seen in the above commands the interfaces 0/1/1, 1/1/1, and 2/1/1 are configured as the backplane interfaces of the three cluster nodes.

2. On the cluster IP address, create VLANs for the backplane interfaces and for the client and server interfaces.

   ```
   //For the backplane interfaces
   add vlan 10
   bind vlan 10 0/1/1 1/1/1 2/1/1

   //For the interfaces that are connected to the client and
   server networks.
   add vlan 20
   bind vlan 20 0/1/2 1/1/2 2/1/2
   ```

3. On the switch, create VLANs for the interfaces corresponding to the backplane interfaces and the client and server interfaces. The following sample configurations are provided for the Cisco C3750 Version 12.2 (40) SE switch. Similar configurations must be performed on other switches.

   ```
   //For the backplane interfaces. Repeat for each interface...
   interface GigabitEthernet1/0/1
   switchport access vlan 100
   switchport mode access
   end

   //For the interfaces connected to the client and server
   networks. Repeat for each interface...
   interface GigabitEthernet1/0/3
   switchport access vlan 200
   switchport mode access
   end
   ```

Common Switch for Client, Server, and Backplane

In this deployment, the client, server, and backplane use dedicated interfaces on the same switch to communicate with the NetScaler cluster.

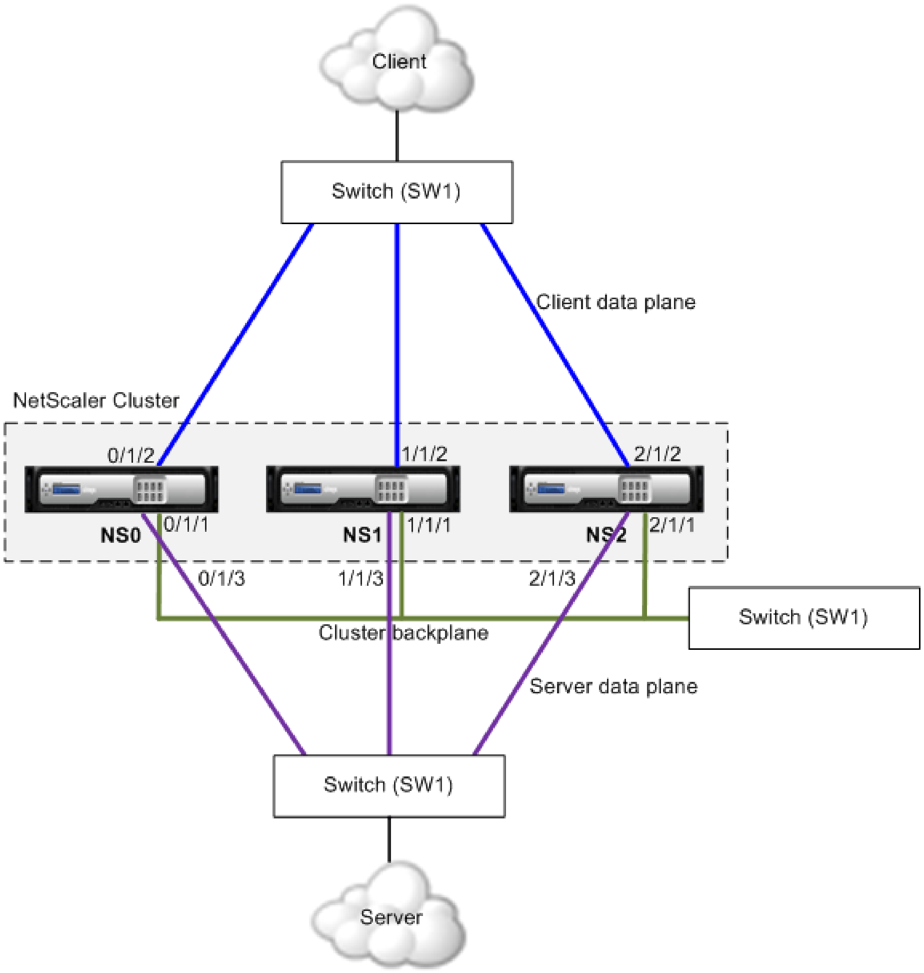

NS0 - nodeId: 0, NSIP: 10.102.29.60

NS1 - nodeId: 1, NSIP: 10.102.29.70

NS2 - nodeId: 2, NSIP: 10.102.29.80

To deploy a cluster with a common switch for the client, server, and backplane

1. Create a cluster of nodes NS0, NS1, and NS2.
   a. Log on to the first node that you want to add to the cluster and do the following:

```
create cluster instance 1
add cluster node 0 10.102.29.60 -state ACTIVE -backplane
0/1/1
enable cluster instance 1
add ns ip 10.102.29.61 255.255.255.255 -type CLIP
save ns config
reboot -warm
```

   b. Log on to the cluster IP address and do the following:

```
add cluster node 1 10.102.29.70 -state ACTIVE -backplane
1/1/1
```

```
add cluster node 2 10.102.29.80 -state ACTIVE -backplane
2/1/1
```

c. Log on to the nodes 10.102.29.70 and 10.102.29.80 to join the nodes to the cluster.

```
join cluster -clip 10.102.29.61 -password nsroot
save ns config
reboot -warm
```

As seen in the above commands the interfaces 0/1/1, 1/1/1, and 2/1/1 are configured as the backplane interfaces of the three cluster nodes.

2. On the cluster IP address, create VLANs for the backplane, client, and server interfaces.

```
//For the backplane interfaces
add vlan 10
bind vlan 10 0/1/1 1/1/1 2/1/1

//For the client-side interfaces
add vlan 20
bind vlan 20 0/1/2 1/1/2 2/1/2

//For the server-side interfaces
add vlan 30
bind vlan 30 0/1/3 1/1/3 2/1/3
```

3. On the switch, create VLANs for the interfaces corresponding to the backplane interfaces and the client and server interfaces. The following sample configurations are provided for the Cisco C3750 Version 12.2 (40) SE switch. Similar configurations must be performed on other switches.

```
//For the backplane interfaces. Repeat for each interface...
interface GigabitEthernet1/0/1
switchport access vlan 100
switchport mode access
end

//For the client interfaces. Repeat for each interface...
interface GigabitEthernet1/0/3
switchport access vlan 200
switchport mode access
end

//For the server interfaces. Repeat for each interface...
interface GigabitEthernet1/0/6
switchport access vlan 300
switchport mode access
end
```

Common Switch for Client and Server and Dedicated Switch for Backplane

In this deployment, the clients and servers use different interfaces on the same switch to communicate with the NetScaler cluster. The cluster backplane uses a dedicated switch for inter-node communication.

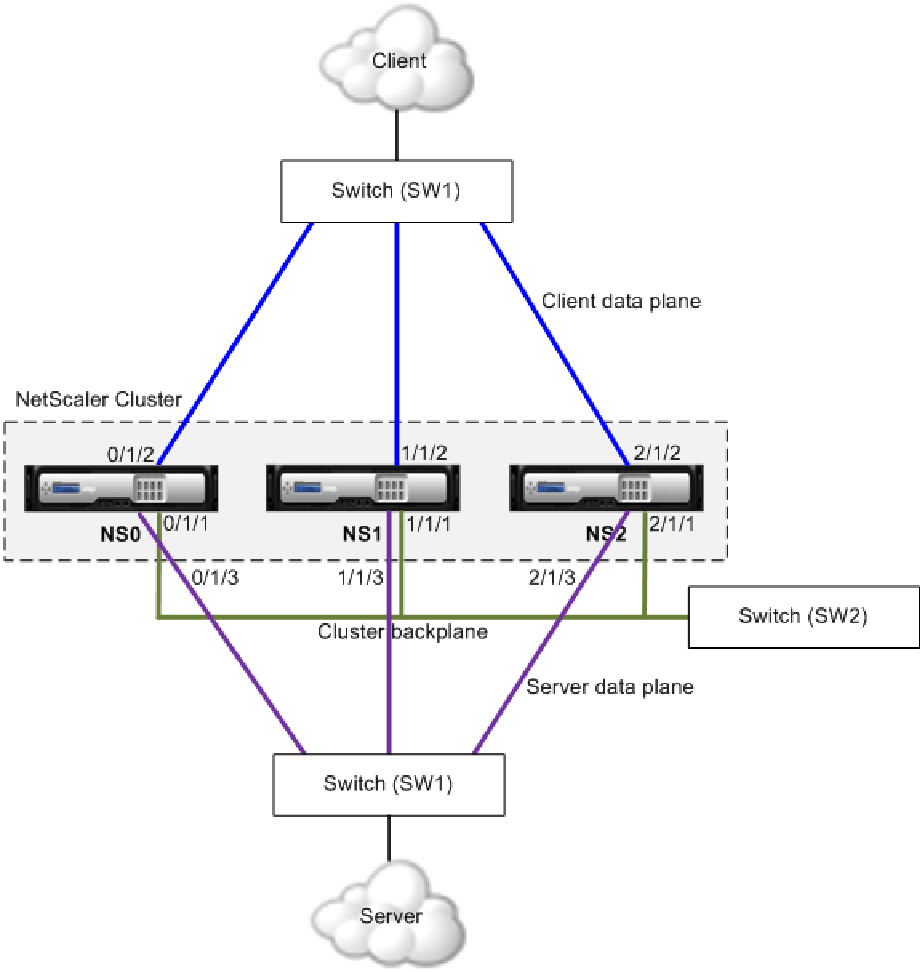

NS0 - nodeId: 0, NSIP: 10.102.29.60

NS1 - nodeId: 1, NSIP: 10.102.29.70

NS2 - nodeId: 2, NSIP: 10.102.29.80

To deploy a cluster with the same switch for the clients and servers and a different switch for the cluster backplane

1. Create a cluster of nodes NS0, NS1, and NS2.

   a. Log on to the first node that you want to add to the cluster and do the following:

```
create cluster instance 1
add cluster node 0 10.102.29.60 -state ACTIVE -backplane
0/1/1
enable cluster instance 1
add ns ip 10.102.29.61 255.255.255.255 -type CLIP
save ns config
reboot -warm
```

b. Log on to the cluster IP address and do the following:

```
add cluster node 1 10.102.29.70 -state ACTIVE -backplane
1/1/1
add cluster node 2 10.102.29.80 -state ACTIVE -backplane
2/1/1
```

c. Log on to the nodes 10.102.29.70 and 10.102.29.80 to join the nodes to the cluster.

```
join cluster -clip 10.102.29.61 -password nsroot
save ns config
reboot -warm
```

As seen in the above commands the interfaces 0/1/1, 1/1/1, and 2/1/1 are configured as the backplane interfaces of the three cluster nodes.

2. On the cluster IP address, create VLANs for the backplane, client, and server interfaces.

```
//For the backplane interfaces
add vlan 10
bind vlan 10 0/1/1 1/1/1 2/1/1

//For the client-side interfaces
add vlan 20
bind vlan 20 0/1/2 1/1/2 2/1/2

//For the server-side interfaces
add vlan 30
bind vlan 30 0/1/3 1/1/3 2/1/3
```

3. On the switch, create VLANs for the interfaces corresponding to the backplane interfaces and the client and server interfaces. The following sample configurations are provided for the Cisco C3750 Version 12.2 (40) SE switch. Similar configurations must be performed on other switches.

```
//For the backplane interfaces. Repeat for each interface...
interface GigabitEthernet1/0/1
switchport access vlan 100
switchport mode access
end

//For the client interfaces. Repeat for each interface...
interface GigabitEthernet1/0/3
switchport access vlan 200
switchport mode access
end

//For the server interfaces. Repeat for each interface...
interface GigabitEthernet1/0/6
switchport access vlan 300
switchport mode access
end
```

Multiple Switches for Each Node

In this deployment, we introduce two client-side switches to ensure redundancy of switches on the client-side. The switches are connected to each other by trunk links. Failure of one switch will not affect the overall working of the cluster.

Note: The same deployment strategy can also be used for server-side connections.

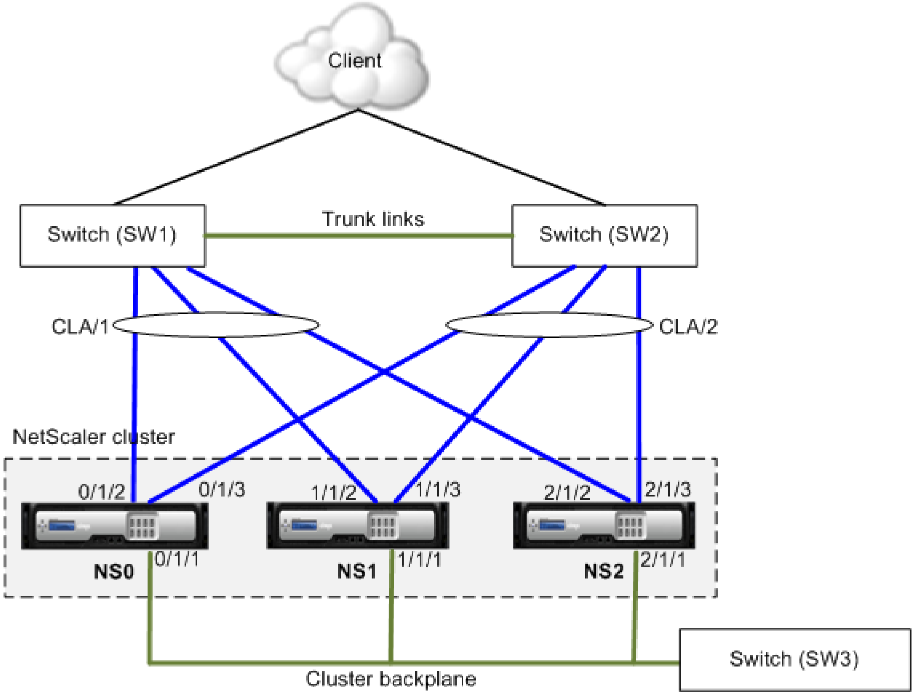

NS0 - nodeId: 0, NSIP: 10.102.29.60

NS1 - nodeId: 1, NSIP: 10.102.29.70

NS2 - nodeId: 2, NSIP: 10.102.29.80

Note: When using trunk links there is possibility of traffic flowing in loops. To avoid this you must make sure the network topology is configured to avoid loops.

To deploy a cluster with each node connected to two switches and the switches connected by trunk links

1. Create a cluster of nodes NS0, NS1, and NS2.
   a. Log on to the first node that you want to add to the cluster and do the following:

```
create cluster instance 1
add cluster node 0 10.102.29.60 -state ACTIVE -backplane
0/1/1
enable cluster instance 1
add ns ip 10.102.29.61 255.255.255.255 -type CLIP
```

```
save ns config
reboot -warm
```

b. Log on to the cluster IP address and do the following:

```
add cluster node 1 10.102.29.70 -state ACTIVE -backplane
1/1/1
add cluster node 2 10.102.29.80 -state ACTIVE -backplane
2/1/1
```

c. Log on to the nodes 10.102.29.70 and 10.102.29.80 to join the nodes to the cluster.

```
join cluster -clip 10.102.29.61 -password nsroot
save ns config
reboot -warm
```

As seen in the above commands the interfaces 0/1/1, 1/1/1, and 2/1/1 are configured as the backplane interfaces of the three cluster nodes.

2. Log on to the cluster IP address and do the following:

a. Create a VLAN for the backplane interfaces.

```
add vlan 10
bind vlan 10 0/1/1 1/1/1 2/1/1
```

b. Create a CLAG for the client-side interfaces with SW1 and SW2.

```
add channel CLA/1 -ifnum 0/1/2 1/1/2 2/1/2 -speed 1000
add channel CLA/2 -ifnum 0/1/3 1/1/3 2/1/3 -speed 1000
```

Different Switch for Every Node

In this deployment, each cluster node is connected to a different switch and trunk links are configured between the switches.

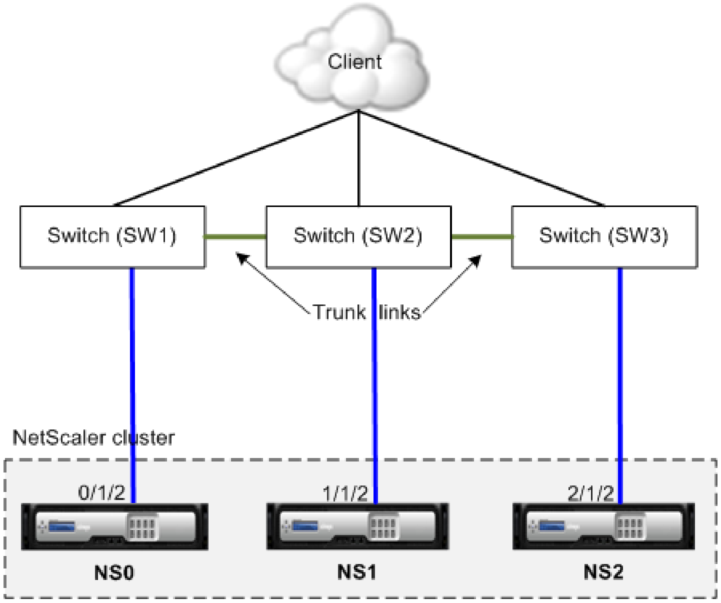

The cluster configurations will be the same as the other deployments scenarios. Most of the client-side configurations will be done on the client-side switches.

Sample Cluster Configurations

The following example can be used to configure a four-node cluster with ECMP, CLAg, or Linksets.

1. Create the cluster.
   a. Log on to first node.
   b. Add the cluster instance.
   ```
   add cluster instance 1
   ```
   c. Add the first node to the cluster.
   ```
   add cluster node 0 10.102.33.184 -backplane 0/1/1
   ```
   d. Enable the cluster instance.
   ```
   enable cluster instance 1
   ```
   e. Add the cluster IP address.
   ```
   add ns ip 10.102.33.185 255.255.255.255 -type CLIP
   ```
   f. Save the configurations.
   ```
   save ns config
   ```
   g. Warm reboot the appliance.
   ```
   reboot -warm
   ```
2. Add the other three nodes to the cluster.
   a. Log on to cluster.
   b. Add the second node to the cluster.
   ```
   add cluster node 1 10.102.33.187 -backplane 1/1/1
   ```
   c. Add the third node to the cluster.
   ```
   add cluster node 2 10.102.33.188 -backplane 2/1/1
   ```
   d. Add the fourth node to the cluster.
   ```
   add cluster node 3 10.102.33.189 -backplane 3/1/1
   ```
3. Join the added nodes to the cluster. This step is not applicable for the first node.
   a. Log on to each newly added node.
   b. Join the node to the cluster.
   ```
   join cluster -clip 10.102.33.185 -password nsroot
   ```
   c. Save the configuration.
   ```
   save ns config
   ```

d. Warm reboot the appliance.

```
reboot -warm
```

4. Configure the NetScaler cluster through the cluster IP address.

```
// Enable load balancing feature
enable ns feature lb

// Add a load balancing virtual server
add lb vserver first_lbvserver http
....
....
```

5. Configure any one of the following (ECMP, Linkset, CLAG) traffic distribution mechanisms for the cluster.

- **ECMP.**

  i. Log on to the cluster.

  ii. Enable the OSPF routing protocol.

  ```
  enable ns feature ospf
  ```

  iii. Add a VLAN.

  ```
  add vlan 97
  ```

  iv. Bind the interfaces of the cluster nodes to the VLAN.

  ```
  bind vlan 97 -ifnum 0/1/4 1/1/4 2/1/4 3/1/4
  ```

  v. Add a spotted SNIP on each node and enable dynamic routing on it.

  ```
  add ns ip 1.1.1.10 255.255.255.0 -ownerNode 0 -
  dynamicRouting ENABLED
  add ns ip 1.1.1.11 255.255.255.0 -ownerNode 1 -
  dynamicRouting ENABLED
  add ns ip 1.1.1.12 255.255.255.0 -ownerNode 2 -
  dynamicRouting ENABLED
  add ns ip 1.1.1.13 255.255.255.0 -ownerNode 3 -
  dynamicRouting ENABLED
  ```

  vi. Bind one of the SNIP addresses to the VLAN.

  ```
  bind vlan 97 -ipAddress 1.1.1.10 255.255.255.0
  ```

  vii. Configure the routing protocol on ZebOS by using vtysh shell.

- **Linksets.** Assume that the node with nodeId 3 is not connected to the switch. You must configure a linkset so that the unconnected node can use the other node interfaces to communicate with the switch.

  i. Log on to the cluster.

  ii. Add a linkset.

  ```
  add linkset LS/1
  ```

  iii. Bind the connected interfaces to the linkset.

  ```
  bind linkset LS/1 -ifnum 0/1/6 1/1/6 2/1/6
  ```

- **Static CLAG.**

i. Log on to the cluster.

ii. Add a CLA channel.

```
add channel CLA/1 -speed 1000
```

iii. Bind the interfaces to the CLA channel.

```
bind channel CLA/1 0/1/5 1/1/5 2/1/5 3/1/5
```

iv. Perform equivalent configuration on the switch.

- **Dynamic CLAG.**

i. Log on to the cluster.

ii. Add the interfaces to the CLA channel.

```
set interface 0/1/5 -lacpmode active -lacpkey 5 -lagtype
cluster
set interface 1/1/5 -lacpmode active -lacpkey 5 -lagtype
cluster
set interface 2/1/5 -lacpmode active -lacpkey 5 -lagtype
cluster
set interface 3/1/5 -lacpmode active -lacpkey 5 -lagtype
cluster
```

iii. Perform equivalent configuration on the switch.

6. Update the state of the cluster nodes to ACTIVE

```
set cluster node 0 -state ACTIVE
set cluster node 1 -state ACTIVE
set cluster node 2 -state ACTIVE
set cluster node 3 -state ACTIVE
```

Troubleshooting the NetScaler Cluster

If a failure occurs in a NetScaler cluster, the first step in troubleshooting is to get information on the cluster instance and the cluster nodes by running the **show cluster instance <clId>** and **show cluster node <nodeId>** commands respectively.

If you are not able to find the issue by using the above two approaches, you can use one of the following:

- **Isolate the source of the failure.** Try bypassing the cluster to reach the server. If the attempt is successful, the problem is probably with the cluster setup.
- **Check the commands recently executed.** Run the **history** command to check the recent configurations performed on the cluster. You can also review the `ns.conf` file to verify the configurations that have been implemented.
- **Check the ns.log files.** Use the log files, available in the `/var/log/` directory of each node, to identify the commands executed, status of commands, and the state changes.
- **Check the newnslog files.** Use the newnslog files, available in the `/var/nslog/` directory of each node, to identify the events that have occurred on the cluster

nodes. You can view multiple newnslog files as a single file, by copying the files to a single directory, and then running the following command:

```
nsconmsg -K newnslog-node<id> -K newnslog.node<id> -d current
```

If you still cannot resolve the issue, you can try tracing the packets on the cluster or use the **show tech support scope cluster** command to send the report to the technical support team.

Tracing the Packets of a NetScaler Cluster

The NetScaler operating system provides a utility called *nstrace* to get a dump of the packets that are received and sent out by an appliance. The utility stores the packets in trace files. You can use these files to debug problems in the flow of packets to the cluster nodes. The trace files must be viewed with the *Wireshark* application. For traces collected in native (.cap) mode, it is important to use the internal version of Wireshark, which can understand native packets.

Some salient aspects of the nstrace utility are:

- Can be configured to trace packets selectively by using classic expressions and default expressions.
- Can capture the trace in multiple formats: nstrace format (.cap) and TCP dump format (.pcap).
- Can aggregate the trace files of all cluster nodes on the CCO.
- Can merge multiple trace files into a single trace file.

You can use the nstrace utility from the NetScaler command line or the NetScaler shell.

To trace packets of a standalone appliance

Run the **start nstrace** command on the appliance. The command creates trace files in the `/var/nstrace/<date-timestamp>` directory. The trace file names are of the form `nstrace<id>.cap`.

You can view the status by executing the **show nstrace** command. You can stop tracing the packets by executing the **stop nstrace** command.

Note: You can also run the nstrace utility from the NetScaler shell by executing the `nstrace.sh` file. Citrix recommends using the nstrace utility through the NetScaler command line.

To trace packets of a cluster

You can trace the packets on all the cluster nodes and obtain all the trace files on the CCO node.

Run the **start nstrace** command on the cluster IP address. The command is propagated and executed on all the cluster nodes. The trace files are stored in individual cluster nodes in the `/var/nstrace/<date-timestamp>` directory. The trace file names are of the form `nstrace<id>_node<id>.cap`.

You can use the trace files of each node to debug the nodes operations. But if you want the trace files of all cluster nodes in one location, you must run the **stop nstrace** command on the cluster IP address. The trace files of all the nodes are downloaded on the cluster configuration coordinator (CCO) node in the `/var/nstrace/<date-timestamp>` directory as follows:

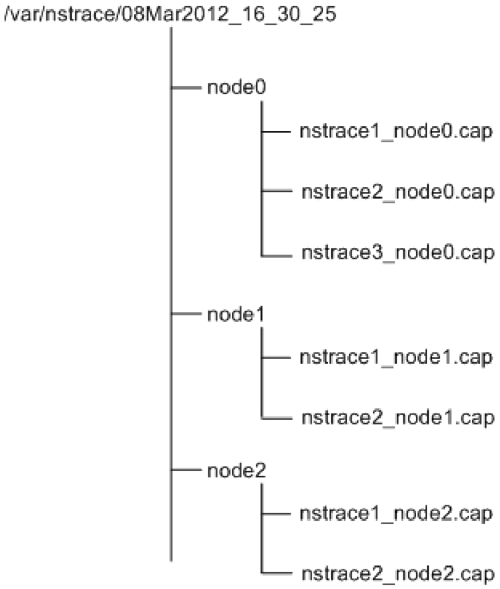

Merge multiple trace files

You can prepare a single file from the trace files obtained from the cluster nodes. The single trace files gives you a cumulative view of the trace of the cluster packets. The trace entries in the single trace file are sorted based on the time the packets were received on the cluster.

To merge the trace files, at the NetScaler shell, type:

nstracemerge.sh -srcdir <DIR> -dstdir <DIR> -filename <name> -filesize <num>

where,

- srcdir is the directory from which the trace files are merged. All trace files within this directory are merged into a single file.
- dstdir is the directory where the merged trace file are created.
- filename is the name of the trace file that is created.
- filesize is the size of the trace file.

Examples

Following are some examples of using the nstrace utility to filter packets.

- To trace the packets on the backplane interfaces of three nodes:

Using classic expressions:

```
start nstrace -filter "INTF == 0/1/1 && INTF == 1/1/1 && INTF
== 2/1/1"
```

Using default expressions:

```
start nstrace -filter "CONNECTION.INTF.EQ("0/1/1") &&
CONNECTION.INTF.EQ("1/1/1") && CONNECTION.INTF.EQ("2/1/1")"
```

- To trace the packets from a source IP address 10.102.34.201 or from a system whose source port is greater than 80 and the service name is not "s1":

  Using classic expressions

  ```
  start nstrace -filter "SOURCEIP == 10.102.34.201 || (SVCNAME !=
  s1 && SOURCEPORT > 80)"
  ```

  Using default expressions

  ```
  start nstrace -filter "CONNECTION.SRCIP.EQ(10.102.34.201) ||
  (CONNECTION.SVCNAME.NE("s1") && CONNECTION.SRCPORT.GT(80))"
  ```

Troubleshooting Common Issues

While joining a node to the cluster, I get the following message, "ERROR: Invalid interface name/number." What must I do to resolve this error?

This error occurs if you provided an invalid or incorrect backplane interface while using the **add cluster node** command to add the node. To resolve this error, verify the interface you provided while adding the node. Make sure that you have not specified the appliance's management interface as the backplane interface, and that the nodeId bit of the interface is the same as the node's Id. For example, if the nodeId is 3, the backplane interface must be 3/<c>/<u>.

While joining a node to the cluster, I get the following message, "ERROR: Clustering cannot be enabled, because the local node is not a member of the cluster." What must I do to resolve this error?

This error occurs when you try to join a node without adding the node's NSIP to the cluster. To resolve this error, you must first add the node's NSIP address to the cluster by using the **add cluster node** command and then execute the **join cluster** command.

While joining a node to the cluster, I get the following message, "ERROR: Connection refused." What must I do to resolve this error?

This error can occur due to the following reasons:

- **Connectivity problems.** The node cannot connect to the cluster IP address. Try pinging the cluster IP address from the node that you are trying to join.
- **Duplicate cluster IP address.** Check to see if the cluster IP address exists on some non-cluster node. If it does, create a new cluster IP address and try re-joining the cluster.

While joining a node to the cluster, I get the following message, "ERROR: License mismatch between CCO and local node." What must I do to resolve this error?

The appliance that you are joining to the cluster must have the same licenses as the CCO. This error occurs when the licenses on the node you are joining do not match the licenses on the CCO. To resolve this error, run the following commands on both the nodes and compare the outputs.

From the command line

- show ns hardware
- show ns license

From the shell

- nsconmsg -g feature -d stats
- ls /nsconfig/license
- View the contents of the /var/log/license.log file

What must I do when the configurations of a cluster node are not in synch with the cluster configurations?

In most cases, the configurations are automatically synchronized between all the cluster nodes. However, if you feel that the configurations are not synchronized on a specific node, you must force the synchronization by executing the **force cluster sync** command from the node that you want to synchronize. For more information, see Synchronizing Cluster Configurations on page 36.

The configurations across the cluster nodes are not synchronized. How can I ensure that the configurations are always in sync?

To ensure synchronization of the cluster configurations across the nodes, run the **save ns config** command after every configuration. Otherwise, the configurations might not be available on the cluster nodes when they reboot.

When configuring a cluster node, I get the following message, "ERROR: Session is read-only; connect to the cluster IP address to modify the configuration."

All configurations on a cluster must be done through the cluster IP address and the configurations are propagated to the other cluster nodes. All sessions established through the NetScaler IP (NSIP) address of individual nodes are read-only.

Why does the node state show "INACTIVE" when the node health shows "UP"?

A healthy node can be in the INACTIVE state for a number of reasons. A scan of `ns.log` or error counters can help you determine the exact reason.

How can I resolve the health of a node when its health shows "Not UP"?

Node health "**Not UP**" indicates that there are some issues with the node. To know the root cause, you must run the **sh cluster node** command. This command displays the node properties and the reason for the node failure.

When I run the set vserver command, I get the following message, "No such resource." What must I do to resolve this issue?

The **set vserver** command is not supported in clustering. The **unset vserver**, **enable vserver**, **disable vserver**, and **rm vserver** commands are also not supported. However, the **show vserver** command is supported.

I cannot configure the cluster over a Telnet session. What must I do?

Over a telnet session, the cluster IP address can be accessed only in read-only mode. Therefore, you cannot configure a cluster over a telnet session.

I notice a significant time difference across the cluster nodes. What must I do to resolve this issue?

When PTP packets are dropped due to backplane switch or if the physical resources are over-committed in a virtual environment, the time will not get synchronized.

To synchronize the times, you must do the following on the cluster IP address:

1. Disable PTP.

   ```
   set ptp -state disable
   ```

2. Configure Network Time Protocol (NTP) for the cluster. For detailed information, see the "Setting Up Clock Synchronization by Using the CLI or the Configuration Utility" chapter of the *Citrix NetScaler Administration Guide*. For a link to the guide, see the Documentation Library on page 65.

FAQs

How many NetScaler appliances can I have in a cluster?

A NetScaler cluster can include as few as 2 or as many as 32 NetScaler nCore hardware or virtual appliances.

I have multiple standalone nodes, each of which has different configurations. Can I add them to a single cluster?

Yes. You can add a maximum of 32 nodes to the cluster. However, the existing configurations of the appliances are cleared when the nodes are added to the cluster. To use the configurations from individual appliances, you must manually prepare a single `*.conf` file of all the configurations, edit the configurations to remove features unsupported by clustering, update the naming convention of interfaces, and then apply the configurations to the CCO by using the **batch** command.

Can I migrate the configurations of a standalone NetScaler appliance or an HA setup to the clustered setup?

No. When a node is added to a clustered setup, the **clear ns config** command (with the extended option) is executed on that appliance. In addition, the SNIP addresses and all VLAN configurations (except default VLAN and NSVLAN) are cleared. Therefore, Citrix recommends backing up the configurations before adding the appliance to a cluster.

Can I automatically detect NetScaler appliances so that I can add them to a cluster?

Yes. The configuration utility allows you to discover NetScaler appliances present in the same subnet as the NSIP address of the CCO. For more information, see Discovering NetScaler Appliances on page 34.

Is the cluster a licensed feature?

Yes, the cluster is a licensed feature. You must have a copy of the cluster license file in the `/nsconfig/license/` directory of all the appliances that you want to add to the cluster. Additionally, all the appliances that you want to add to the cluster, must also have the same license files available.

Can a NetScaler appliance be a part of multiple clusters?

No. An appliance can belong to only one cluster.

Can a node that is not connected to the client or server network still serve traffic?

Yes. The NetScaler cluster supports a traffic distribution mechanism called linksets, which allows unconnected nodes to serve traffic by using the interfaces of connected nodes. The unconnected nodes communicate with the connected nodes through the cluster backplane.

What will happen if the cluster license on a node has expired?

If the cluster license on a node expires when the node is running, the cluster is not affected. However, when you reboot that node, the cluster is operationally disabled on this node, therefore the node will not be able to serve traffic. To rectify the issue and make the node active, you must upload a new license and warm reboot the appliance.

Why are the network interfaces of a cluster represented using a 3-tuple (n/u/c) notation instead of the regular 2-tuple (u/c) notation?

When an appliance is part of a cluster, you must be able to identify the node to which the network interface belongs. So, the network interface naming convention for cluster nodes is modified from u/c to n/u/c, where n denotes the node Id.

What is a striped IP address?

Any IP address (VIP or SNIP) that is defined on the cluster is, by default, a striped IP address. Striped IP addresses are active on all nodes of the cluster.

What is a spotted IP address? Can I change the ownership of a spotted IP address at run time?

A spotted IP address is an IP address that is active and owned exclusively by one node of the cluster. Spotted IP address must be defined through the cluster IP address, by specifying the owner node in the **add ns ip** command.

You cannot change the ownership of a spotted IP address at run time. To change the ownership, you must first delete the IP address and add it again by specifying the new owner.

What is CCO?

CCO is the abbreviated form of *Configuration Coordinator*. This node owns the cluster IP address and coordinates all the cluster configurations.

What is a cluster IP address? What is its subnet mask?

The cluster IP address is the management address of a NetScaler cluster. All cluster configurations must be performed by accessing the cluster through this address. The subnet mask of the cluster IP address is fixed at 255.255.255.255.

When I added the first node to the cluster it was the configuration coordinator (CCO). Now, another node is displayed as the CCO. Why?

When a cluster is created, the first node becomes the CCO. The cluster IP address is owned by that node. However, the CCO is not a fixed node. It can change over time due to various reasons. In that case, the cluster elects a new CCO and assigns the cluster IP address to the new CCO.

Can I execute commands from the NSIP address of a cluster node?

No. Access to individual cluster nodes through the NetScaler IP (NSIP) addresses is read-only. This means that when you log on to the NSIP address of a cluster node you can only view the configurations and the statistics. You cannot perform any configurations. However, there are some operations you can execute from the NSIP address of a cluster node. For more information, see Operations Supported on Individual Nodes on page 63.

Can I disable configuration propagation among cluster nodes?

No, you cannot explicitly disable the propagation of cluster configurations among cluster nodes. However, configuration propagation can get disabled automatically during software upgrade or downgrade on account of version mismatch.

How can I delete a cluster and all the nodes of the cluster?

To delete a cluster and all the nodes of the cluster, you must remove each node individually as described in Removing a Cluster Node on page 22.

Can I change the NSIP address or change the NSVLAN of a NetScaler appliance when it is a part of the cluster?

No. To make such changes you must first remove the appliance from the cluster, perform the changes, and then add the appliance to the cluster.

Does the NetScaler cluster support L2 and L3 Virtual Local Area Networks (VLANs)?

Yes, a NetScaler cluster supports VLANs between cluster nodes. The VLANs must be configured on the cluster IP address.

- **L2 VLAN.** You can create a layer2 VLAN by binding interfaces that belong to different nodes of the cluster.
- **L3 VLAN.** You can create a layer3 VLAN by binding IP addresses that belong to different nodes of the cluster. The IP addresses must belong to the same subnet. Make sure that one of the following criteria is satisfied. Otherwise, the L3 VLAN bindings can fail:
  - All nodes have an IP address on the same subnet as the one bound to the VLAN.
  - The cluster has a striped IP address and the subnet of that IP address is bound to the VLAN.

  When you add a new node to a cluster that has only spotted IPs, the sync happens before spotted IP addresses are assigned to that node. In such cases, L3 VLAN bindings can be lost. To avoid this loss, either add a striped IP or add the L3 VLAN bindings on the NSIP of the newly added node.

Why are the VLANs and VLAN bindings deleted when a NetScaler appliance is added to the cluster?

When a NetScaler appliance is added to a clustered setup, the **clear ns config** command (with the extended option) is executed on that appliance. In addition, the SNIP addresses and all VLAN configurations (except the default VLAN and NSVLAN) are cleared.

How can I configure SNMP on a NetScaler cluster?

SNMP monitors the cluster, and all the nodes of the cluster, in the same way as it monitors a standalone NetScaler appliance. The only difference is that SNMP on a cluster must be configured through the cluster IP address. When generating hardware specific traps, two additional varbinds are included to identify the node of the cluster: node ID and NSIP of the node.

For detailed information about configuring SNMP, see the "SNMP" chapter of the *Citrix NetScaler Administration Guide*. For a link to the guide, see the Documentation Library on page 65.

What details must I have available when I contact technical support for cluster-related issues?

The NetScaler provides a **show techsupport -scope cluster** command that extracts configuration data, statistical information, and logs of all the cluster nodes. You must run this command on the cluster IP address.

The output of this command is saved in a file named `collector_cluster_<nsip_CCO>_P_<date-timestamp>.tar.gz` which is available in the `/var/tmp/support/cluster/` directory of the CCO.

Send this archive to the technical support team to debug the issue.

Operations Supported on Individual Nodes

All cluster configurations are performed on the CCO through the cluster IP address and these configurations are propagated to the cluster nodes. However, there are some operations that can be performed on individual cluster nodes by accessing them through their NetScaler IP (NSIP) addresses.

- enable cluster instance
- disable cluster instance
- set cluster instance
- rm cluster instance
- set cluster node
- rm cluster node
- force cluster sync
- sync cluster files
- send arp all
- start nstrace
- stop nstrace
- show nstrace
- set interface
- enable interface
- disable interface
- save ns config
- reboot

Note: All the show and stat commands are allowed as they do not involve any change in configurations.

Appendix A

Documentation Library

Topics:

- *Release Notes*
- *Quick Start Guides*
- *Configuration Guides*
- *Reference Guides*

This appendix contains links to various NetScaler guides. You can either click the respective document ID to open the PDF version of the guide, or use the ID to search the guide in the Citrix Knowledge Center website available at http://support.citrix.com.

To search the guide on Citrix Knowledge Center website

1. Open the http://support.citrix.com link in a web browser.
2. Type the document ID in the Knowledge Center search text box and click **Search**.
3. Select the appropriate link from the search results.

Release Notes

Title	Document ID
Citrix NetScaler Release Notes	CTX132356

Quick Start Guides

Title	Document ID
Citrix NetScaler Quick Start Guide for NetScaler MPX	CTX132374
Citrix NetScaler Quick Start Guide for NetScaler MPX 5500	CTX132371
Citrix NetScaler Quick Start Guide for NetScaler MPX 7500, 9500	CTX132370
Citrix NetScaler Quick Start Guide for NetScaler MPX 9700, 10500, 12500, 15500	CTX132373
Citrix NetScaler Quick Start Guide for MPX 11500, 13500, 14500, 16500, 18500	CTX132379
Citrix NetScaler Quick Start Guide MPX 17550/19550/20550/21550	CTX132380
Citrix NetScaler Quick Start Guide for NetScaler MPX 17500, 19500, 21500	CTX132377
Citrix NetScaler Quick Start Guide for SDX 11500, 13500, 14500, 16500, 18500, 20500	CTX132785
Citrix NetScaler Quick Start Guide for SDX 17500/19500/21500	CTX132784
Citrix NetScaler Quick Start Guide for SDX 17550/19550/20550/21550	CTX132783

Configuration Guides

Title	Document ID
Citrix NetScaler Administration Guide	CTX132357
Citrix NetScaler AppExpert Guide	CTX132358
Citrix NetScaler Application Optimization Guide	CTX132361
Citrix NetScaler Application Security Guide	CTX132366
Citrix NetScaler Clustering Guide	CTX132840
Citrix Application Firewall Guide	CTX132360
Citrix NetScaler Getting Started Guide	CTX132368
Citrix Hardware Installation and Setup Guide	CTX132365
Citrix NetScaler Migration Guide	CTX132364
Citrix NetScaler Networking Guide	CTX132369
Citrix NetScaler Policy Configuration and Reference Guide	CTX132362
Citrix NetScaler SDX Administration	CTX132782
Citrix NetScaler Traffic Management Guide	CTX132359
Citrix NetScaler VPX Getting Started Guide	CTX132363

Reference Guides

Title	Document ID
Citrix NetScaler Command Reference Guide	CTX132384
Citrix NetScaler Developers Guide	CTX132367
Citrix NetScaler Glossary	CTX132383
Citrix NetScaler Log Message Reference	CTX132382
Citrix NetScaler SNMP OID Reference	CTX132381