



CHAPTER 18

Configuring GLBP

This chapter describes how to configure the Gateway Load Balancing Protocol (GLBP) on the Cisco NX-OS device.

This chapter includes the following sections:

- [Information About GLBP, page 18-1](#)
- [Licensing Requirements for GLBP, page 18-6](#)
- [Prerequisites for GLBP, page 18-7](#)
- [Guidelines and Limitations for GLBP, page 18-7](#)
- [Default Settings, page 18-7](#)
- [Configuring GLBP, page 18-8](#)
- [Verifying the GLBP Configuration, page 18-17](#)
- [Configuration Examples for GLBP, page 18-17](#)
- [Additional References, page 18-17](#)
- [Feature History for GLBP, page 18-18](#)

Information About GLBP

GLBP provides path redundancy for IP by sharing protocol and Media Access Control (MAC) addresses between redundant gateways. Additionally, GLBP allows a group of Layer 3 routers to share the load of the default gateway on a LAN. A GLBP router can automatically assume the forwarding function of another router in the group if the other router fails.

This section includes the following topics:

- [GLBP Overview, page 18-2](#)
- [GLBP Active Virtual Gateway, page 18-2](#)
- [GLBP Virtual MAC Address Assignment, page 18-2](#)
- [GLBP Virtual Gateway Redundancy, page 18-3](#)
- [GLBP Virtual Forwarder Redundancy, page 18-3](#)
- [GLBP Authentication, page 18-4](#)
- [GLBP Load Balancing and Tracking, page 18-5](#)
- [High Availability and Extended Nonstop Forwarding, page 18-6](#)

Send document comments to nexus7k-docfeedback@cisco.com.

- [Virtualization Support, page 18-6](#)

GLBP Overview

GLBP provides automatic [gateway](#) backup for IP hosts configured with a single default gateway on an IEEE 802.3 LAN. Multiple routers on the LAN combine to offer a single virtual first-hop IP gateway while sharing the IP packet forwarding load. Other routers on the LAN might act as redundant GLBP gateways that become active if any of the existing forwarding gateways fail.

GLBP performs a similar function to the Hot Standby Redundancy Protocol (HSRP) and the Virtual Router Redundancy Protocol (VRRP). HSRP and VRRP allow multiple routers to participate in a virtual group configured with a virtual IP address. These protocols elect one member as the active router to forward packets to the virtual IP address for the group. The other routers in the group are redundant until the active router fails.

GLBP performs an additional load balancing function that the other protocols do not provide. GLBP load balances over multiple routers (gateways) using a single virtual IP address and multiple virtual MAC addresses. GLBP shares the forwarding load among all routers in a GLBP group instead of allowing a single router to handle the whole load while the other routers remain idle. You configure each host with the same virtual IP address, and all routers in the virtual group participate in forwarding packets. GLBP members communicate between each other using periodic hello messages.

GLBP Active Virtual Gateway

GLBP prioritizes gateways to elect an active virtual gateway ([AVG](#)). If multiple gateways have the same priority, the gateway with the highest real IP address becomes the AVG. The AVG assigns a virtual MAC address to each member of the GLBP group. Each member is the active virtual forwarder ([AVF](#)) for its assigned virtual MAC address, forwarding packets sent to its assigned virtual MAC address.

The AVG also answers Address Resolution Protocol (ARP) requests for the virtual IP address. Load sharing is achieved when the AVG replies to the ARP requests with different virtual MAC addresses.



Note

Packets received on a routed port destined for the GLBP virtual IP address terminate on the local router, regardless of whether that router is the active GLBP router or a redundant GLBP router. This termination includes ping and Telnet traffic. Packets received on a Layer 2 (VLAN) interface destined for the GLBP virtual IP address terminate on the active router.

GLBP Virtual MAC Address Assignment

The AVG assigns the virtual MAC addresses to each member of the group. The group members request a virtual MAC address after they discover the AVG through hello messages. The AVG assigns the next MAC address based on the load-balancing algorithm selected (see the [“GLBP Load Balancing and Tracking” section on page 18-5](#)). A gateway that is assigned with a virtual MAC address by the AVG is the primary virtual forwarder. The other members of the GLBP group that learn the virtual MAC addresses from hello messages are secondary virtual forwarders.

Send document comments to nexus7k-docfeedback@cisco.com.

GLBP Virtual Gateway Redundancy

GLBP provides virtual gateway redundancy. A member in a group can be in the active, standby, or listen state. GLBP uses a priority algorithm to elect one gateway as the AVG and elect another gateway as the standby virtual gateway. The remaining gateways go into the listen state. You can configure the GLBP priority on each gateway. If the GLBP priority is identical on multiple gateways, GLBP uses the gateway with the highest IP address as the AVG.

If an AVG fails, the standby virtual gateway assumes responsibility for the virtual IP address. GLBP elects a new standby virtual gateway from the gateways in the listen state.

GLBP Virtual Forwarder Redundancy

GLBP provides virtual forwarder redundancy. Virtual forwarder redundancy is similar to virtual gateway redundancy with an active virtual forwarder (AVF). If the AVF fails, a secondary virtual forwarder in the listen state assumes responsibility for the virtual MAC address. This secondary virtual forwarder is also a primary virtual forwarder for a different virtual MAC address. GLBP migrates hosts away from the old virtual MAC address of the failed AVF, using the following two timers:


- Redirect timer—Specifies the interval during which the AVG continues to redirect hosts to the old virtual MAC address. When the redirect time expires, the AVG stops using the old virtual MAC address in ARP replies, although the secondary virtual forwarder continues to forward packets that were sent to the old virtual MAC address.
- Secondary hold timer—Specifies the interval during which the virtual MAC address is valid. When the secondary hold time expires, GLBP removes the virtual MAC address from all gateways in the GLBP group and load balances the traffic over the remaining AVFs. The expired virtual MAC address becomes eligible for reassignment by the AVG.

GLBP uses hello messages to communicate the current state of the timers.

In [Figure 18-1](#), router A is the AVG for a GLBP group and is responsible for the virtual IP address 192.0.2.1. Router A is also an AVF for the virtual MAC address 0007.b400.0101. Router B is a member of the same GLBP group and is designated as the AVF for the virtual MAC address 0007.b400.0102. Client 1 has a default gateway IP address of 192.0.2.1, the virtual IP address, and a gateway MAC address of 0007.b400.0101 that points to router A. Client 2 shares the same default gateway IP address but receives the gateway MAC address 0007.b400.0102 because router B is sharing the traffic load with router A.

Send document comments to nexus7k-docfeedback@cisco.com.

Figure 18-1 GLBP Topology



If router A becomes unavailable, client 1 does not lose access to the WAN because router B assumes responsibility for forwarding packets sent to the virtual MAC address of router A and for responding to packets sent to its own virtual MAC address. Router B also assumes the role of the AVG for the entire GLBP group. Communication for the GLBP members continues despite the failure of a router in the GLBP group.

GLBP Authentication

GLBP has three authentication types:

- MD5 authentication
- Plain text authentication
- No authentication

MD5 authentication provides greater security than plain text authentication. MD5 authentication allows each GLBP group member to use a secret key to generate a keyed MD5 hash that is part of the outgoing packet. At the receiving end, a keyed hash of an incoming packet is generated. If the hash within the incoming packet does not match the generated hash, the packet is ignored. The key for the MD5 hash can either be given directly in the configuration using a key string or supplied indirectly through a key chain.

You can also choose to use a simple password in plain text to authenticate GLBP packets, or choose no authentication for GLBP.

GLBP rejects packets in any of the following cases:

- The authentication schemes differ on the router and in the incoming packet.
- MD5 digests differ on the router and in the incoming packet.
- Text authentication strings differ on the router and in the incoming packet.

Send document comments to nexus7k-docfeedback@cisco.com.

GLBP Load Balancing and Tracking

You can configure the following load-balancing methods for GLBP:

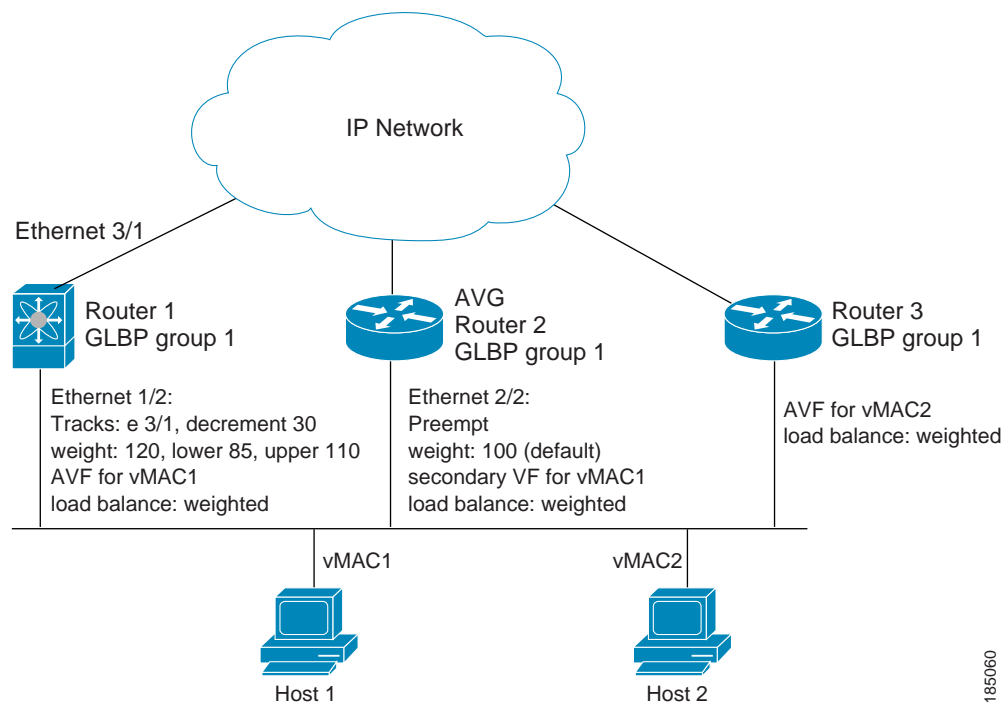
- Round-robin—GLBP cycles through the virtual MAC addresses sent in ARP replies, load balancing the traffic across all the AVFs.
- Weighted—AVG uses the advertised weight for an AVF to decide the load directed to the AVF. A higher weight means that the AVG directs more traffic to the AVF.
- Host dependent—GLBP uses the MAC address of the host to determine which virtual MAC address to direct the host to use. This algorithm guarantees that a host gets the same virtual MAC address if the number of virtual forwarders does not change.

The default for IPv4 networks is round-robin. You can disable all load balancing for GLBP on an interface. If you do not configure load balancing, the AVG handles all traffic for the hosts while the other GLBP group members are in standby or listen mode.

You can configure GLBP to track an interface or routes and enable the secondary virtual forwarder to take over if a tracked link goes down. GLBP tracking uses weighted load-balancing to determine whether a GLBP group member acts as an AVF. You must configure the initial weighting values and optional thresholds to enable or disable this group member as an AVF. You can also configure the interface to track and the value that reduces the interface's weighting if the interface goes down. When the GLBP group weighting drops below the lower threshold, the member is no longer an AVF and a secondary virtual forwarder takes over. When the weighting rises above the upper threshold, the member can resume its role as an AVF.

Figure 18-2 shows an example of GLBP tracking and weighting.

Figure 18-2 GLBP Object Tracking and Weighting



185060

Send document comments to nexus7k-docfeedback@cisco.com.

In [Figure 18-2](#), the Ethernet 1/2 interface on router 1 is the gateway for host 1 (the AVF for virtual MAC address, vMAC1), while Ethernet 2/2 on router 2 acts as a secondary virtual forwarder for Host 1. Ethernet 1/2 tracks Ethernet 3/1, which is the network connection for router 1. If Ethernet 3/1 goes down, the weighting for Ethernet 1/2 drops to 90. Ethernet 2/2 on router 2 preempts Ethernet 1/2 and takes over as AVF because it has the default weighting of 100 and is configured to preempt the AVF.

See the [“Configuring GLBP Weighting and Tracking” section on page 18-11](#) for details about configuring weighting and tracking.

High Availability and Extended Nonstop Forwarding

GLBP supports high availability through stateful restarts and stateful switchovers. A stateful restart occurs when the GLBP process fails and is restarted. A stateful switchover occurs when the active supervisor switches to the standby supervisor. Cisco NX-OS applies the run-time configuration after the switchover.

If GLBP hold timers are configured for short time periods, these timers might expire during a controlled switchover or in-service software upgrade (ISSU). GLBP supports extended non-stop forwarding (NSF) to temporarily extend these GLBP hold timers during a controlled switchover or ISSU.

With extended NSF configured, GLBP sends hello messages with the extended timers. GLBP peers update their hold timers with these new values. The extended timers prevent unnecessary GLBP state changes during the switchover or ISSU. After the switchover or ISSU event, GLBP restores the hold timers to their original configured values. If the switchover fails, GLBP restores the hold timers after the extended hold timer values expire.

See the [“Configuring Extended Hold Timers for GLBP” section on page 18-15](#) for more information.

Virtualization Support

GLBP supports virtual routing and forwarding (VRF) instances. VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF.

If you change the VRF membership of an interface, Cisco NX-OS removes all Layer 3 configuration, including GLBP.

For more information, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 5.x*, and see [Chapter 14, “Configuring Layer 3 Virtualization.”](#)

Licensing Requirements for GLBP

The following table shows the licensing requirements for this feature:

| Product | License Requirement |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cisco NX-OS | GLBP requires no license. Any feature not included in a license package is bundled with the Cisco NX-OS system images and is provided at no extra charge to you. For a complete explanation of the Cisco NX-OS licensing scheme, see the <i>Cisco NX-OS Licensing Guide</i> . |

Send document comments to nexus7k-docfeedback@cisco.com.

Prerequisites for GLBP

GLBP has the following prerequisites:

- Globally enable the GLBP feature (see the “[Enabling GLBP](#)” section on page 18-8).
- You can only configure GLBP on Layer 3 interfaces (see the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 5.x*, and the *Cisco DCNM Interfaces Configuration Guide, Release 5.x*).
- If you configure VDCs, install the Advanced Services license and enter the desired VDC (see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 5.x*).

Guidelines and Limitations for GLBP

GLBP has the following configuration guidelines and limitations:

- You should configure all customization options for GLBP on all GLBP member gateways before enabling a GLBP group by configuring a virtual IP address.
- You must configure an IP address for the interface that you configure GLBP on and enable that interface before GLBP becomes active.
- The GLBP virtual IP address must be in the same subnet as the interface IP address.
- We recommend that you do not configure more than one first-hop redundancy protocol on the same interface.
- Cisco NX-OS removes all Layer 3 configuration on an interface when you change the VDC, interface VRF membership, port channel membership, or when you change the port mode to Layer 2.
- Cisco NX-OS does not support GLBP group configuration on interface secondary subnets.
- Cisco NX-OS does not support GLBP for IPv6.

Default Settings

[Table 18-1](#) lists the default settings for GLBP parameters.

Table 18-1 *Default GLBP Parameters*

| Parameters | Default |
|----------------------------|-------------------|
| Authentication | No authentication |
| Extended hold timer | 10 seconds |
| Forwarder preemption delay | 30 seconds |
| Forwarder timeout | 14400 seconds |
| Hello timer | 3 seconds |
| Hold timer | 10 seconds |
| GLBP feature | Disabled |
| Load balancing | Round robin |

Send document comments to nexus7k-docfeedback@cisco.com.

Table 18-1 *Default GLBP Parameters (continued)*

| Parameters | Default |
|----------------|-------------|
| Preemption | Disabled |
| Priority | 100 |
| Redirect timer | 600 seconds |
| Weighting | 100 |

Configuring GLBP

This section includes the following topics:

- [Enabling GLBP, page 18-8](#)
- [Configuring GLBP Authentication, page 18-9](#)
- [Configuring GLBP Load Balancing, page 18-11](#)
- [Configuring GLBP Weighting and Tracking, page 18-11](#)
- [Customizing GLBP, page 18-14](#)
- [Configuring Extended Hold Timers for GLBP, page 18-15](#)
- [Enabling a GLBP Group, page 18-15](#)



Note

If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Enabling GLBP

You must enable GLBP before you can configure and enable any GLBP groups.

BEFORE YOU BEGIN

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

DETAILED STEPS

To enable GLBP, use the following command in global configuration mode:

| Command | Purpose |
|-------------------------------------------------|---------------|
| feature glbp | Enables GLBP. |
| Example: switch(config)# feature glbp | |

Send document comments to nexus7k-docfeedback@cisco.com.

To disable GLBP in a VDC and remove all associated configuration, use the following command in global configuration mode:

| Command | Purpose |
|----------------------------------------------------|-------------------------|
| no feature glbp | Disables GLBP in a VDC. |
| Example: switch(config)# no feature glbp | |

Configuring GLBP Authentication

You can configure GLBP to authenticate the protocol using cleartext or an MD5 digest. MD5 authentication uses a key chain (see the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide, Release 5.x*).

BEFORE YOU BEGIN

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Enable GLBP (see the “[Enabling GLBP](#)” section on page 18-8).



Note

You must configure the same authentication and keys on all members of the GLBP group.

SUMMARY STEPS

1. **configure terminal**
2. **interface** *interface-type slot/port*
3. **ip** *ip-address/length*
4. **glbp** *group-number*
5. **authentication text** *string*
or
authentication md5 {**key-chain** *key-chain* | **key-string** {*text* / **encrypted** *text*}
6. **ip** [*ip-address* [**secondary**]]
7. (Optional) **show glbp** [**group** *group-number*]
8. (Optional) **copy running-config startup-config**

[Send document comments to nexus7k-docfeedback@cisco.com.](mailto:nexus7k-docfeedback@cisco.com)

DETAILED STEPS

| | Command | Purpose |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | configure terminal Example: switch# configure terminal switch(config)# | Enters configuration mode. |
| Step 2 | interface interface-type slot/port Example: switch(config)# interface ethernet 1/2 switch(config-if)# | Enters interface configuration mode. |
| Step 3 | ip ip-address/length Example: switch(config-if)# ip 192.0.2.1/8 | Configures the IPv4 address for the interface. |
| Step 4 | glbp group-number Example: switch(config-if)# glbp 1 switch(config-if-glbp)# | Creates a GLBP group and enters GLBP configuration mode. The range is from 0 to 1024. |
| Step 5 | authentication text string Example: switch(config-if-glbp)# authentication text mypassword | Configures cleartext authentication for GLBP on this interface. |
| | authentication md5 {key-chain key-chain key-string {text encrypted text}} Example: switch(config-if-glbp)# authentication md5 key-chain glbp-keys | Configures MD5 authentication for GLBP on this interface. |
| Step 6 | ip [ip-address [secondary]] Example: switch(config-if-glbp)# ip 192.0.2.10 | Enables GLBP on an interface and identifies the primary IP address of the virtual gateway. After you identify a primary IP address, you can use the glbp group ip command again with the secondary keyword to indicate additional IP addresses supported by this group. If you only use the ip keyword, GLBP learns the virtual IP address from the neighbors. |
| Step 7 | show glbp [group group-number] Example: switch(config-if-glbp)# show glbp 1 | (Optional) Displays GLBP information. |
| Step 8 | copy running-config startup-config Example: switch(config-if-glbp)# copy running-config startup-config | (Optional) Saves this configuration change. |

Send document comments to nexus7k-docfeedback@cisco.com.

This example shows how to configure MD5 authentication for GLBP on Ethernet 1/2 after creating the key chain:

```
switch# configure terminal
switch(config)# key chain glbp-keys
switch(config-keychain)# key 0
switch(config-keychain-key)# key-string 7 zqdest
switch(config-keychain-key) accept-lifetime 00:00:00 Jun 01 2008 23:59:59 Sep 12 2008
switch(config-keychain-key) send-lifetime 00:00:00 Jun 01 2008 23:59:59 Aug 12 2008
switch(config-keychain-key) key 1
switch(config-keychain-key) key-string 7 uaeqdyito
switch(config-keychain-key) accept-lifetime 00:00:00 Aug 12 2008 23:59:59 Dec 12 2008
switch(config-keychain-key) send-lifetime 00:00:00 Sep 12 2008 23:59:59 Nov 12 2008
switch(config)# interface ethernet 1/2
switch(config-if)# glbp 1
switch(config-if-glbp)# authenticate md5 key-chain glbp-keys
switch(config-if-glbp)# copy running-config startup-config
```

Configuring GLBP Load Balancing

You can configure GLBP to use load balancing based on round-robin, weighted, or host-dependent methods (see the “[GLBP Load Balancing and Tracking](#)” section on page 18-5).

To configure GLBP load balancing, use the following command in GLBP configuration mode:

| Command | Purpose |
|---------------------------------------------------------------------------------|------------------------------------------------------------------|
| <code>load-balancing [host-dependent round-robin weighted]</code> | Sets the GLBP load-balancing method. The default is round-robin. |
| Example: <code>switch(config-if-glbp)# load-balancing weighted</code> | |

Configuring GLBP Weighting and Tracking

You can configure GLBP weighting values and object tracking to work with the GLBP weighted load-balancing method.

You can optionally configure the interface to preempt an AVF if the interface was originally assigned with the virtual MAC address or if this interface has a higher weight than the AVF.

BEFORE YOU BEGIN

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Enable GLBP (see the “[Enabling GLBP](#)” section on page 18-8).

SUMMARY STEPS

1. `configure terminal`
2. `track object-id interface interface-type number {ip routing | line-protocol}`
`track object-id ip route ip-prefix/length reachability`
3. `interface interface-type slot/port`

[Send document comments to nexus7k-docfeedback@cisco.com.](mailto:nexus7k-docfeedback@cisco.com)

4. **ip** *ip-address/length*
5. **glbp** *group-number*
6. **weighting** *maximum* [**lower** *lower*] [**upper** *upper*]
7. **weighting track** *object-number* [**decrement** *value*]
8. (Optional) **forwarder preempt** [**delay** *minimum* *seconds*]
9. **ip** [*ip-address* [**secondary**]]
10. (Optional) **show glbp** *interface-type number*
11. (Optional) **copy running-config startup-config**

DETAILED STEPS

| | Command | Purpose |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | configure terminal Example: switch# configure terminal switch(config)# | Enters configuration mode. |
| Step 2 | track <i>object-id interface</i> <i>interface-type number</i> { ip routing line-protocol } Example: switch(config)# track 1 interface ethernet 2/2 line-protocol switch(config-track# | Configures the interface that this GLBP interface tracks. Changes in the state of the interface affect the priority of this GLBP interface as follows: <ul style="list-style-type: none"> • You configure the interface and corresponding object number that you use with the track command in GLBP configuration mode. • The line-protocol keyword tracks whether the interface is up. The ip keyword also checks that IP routing is enabled on the interface and an IP address is configured. |
| | track <i>object-id ip route</i> <i>ip-prefix/length</i> reachability Example: switch(config)# track 2 ip route 192.0.2.0/8 reachability switch(config-track# | Creates a tracked object for a route and enters tracking configuration mode. The <i>object-id</i> range is from 1 to 500. |
| Step 3 | interface <i>interface-type slot/port</i> Example: switch(config)# interface ethernet 1/2 switch(config-if)# | Enters interface configuration mode. |
| Step 4 | ip <i>ip-address/length</i> Example: switch(config-if)# ip 192.0.2.1/8 | Configures the IPv4 address for the interface. |
| Step 5 | glbp <i>group-number</i> Example: switch(config-if)# glbp 1 switch(config-if-glbp)# | Creates a GLBP group and enters GLBP configuration mode. |

Send document comments to nexus7k-docfeedback@cisco.com.

| | Command | Purpose |
|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 6 | weighting <i>maximum</i> [lower <i>lower</i>] [upper <i>upper</i>] Example: switch(config-if-glbp)# weighting 110 lower 95 upper 105 | Specifies the initial weighting value and the upper and lower thresholds for a GLBP gateway. The maximum range is from 1 to 254. The default weighting value is 100. The lower range is from 1 to 253. The upper range is from 1 to 254. |
| Step 7 | weighting track <i>object-number</i> [decrement <i>value</i>] Example: switch(config-if-glbp)# weighting track 2 decrement 20 | Specifies an object to be tracked that affects the weighting of a GLBP gateway. The <i>value</i> argument specifies a reduction in the weighting of a GLBP gateway when a tracked object fails. The range is from 1 to 255. |
| Step 8 | forwarder preempt [delay <i>minimum</i> <i>seconds</i>] Example: switch(config-if-glbp)# forwarder preempt delay minimum 60 | (Optional) Configures the router to take over as AVF for a GLBP group if the current AVF for a GLBP group falls below its low weighting threshold. The range is from 0 to 3600 seconds. This command is enabled by default with a delay of 30 seconds. |
| Step 9 | ip [<i>ip-address</i> [secondary]] Example: switch(config-if-glbp)# ip 192.0.2.10 | Enables GLBP on an interface and identifies the primary IP address of the virtual gateway. After you identify a primary IP address, you can use the glbp group ip command again with the secondary keyword to indicate additional IP addresses supported by this group. If you only use the ip keyword, GLBP learns the virtual IP address from the neighbors. |
| Step 10 | show glbp <i>interface-type number</i> Example: switch(config-if-glbp)# show glbp ethernet 1/2 | (Optional) Displays GLBP information for an interface. |
| Step 11 | copy running-config startup-config Example: switch(config-if-glbp)# copy running-config startup-config | (Optional) Saves this configuration change. |

The following example shows how to configure GLBP weighting and tracking on Ethernet 1/2:

```
switch# configure terminal
switch(config)# track 2 interface ethernet 2/2 ip routing
switch(config)# interface ethernet 1/2
switch(config-if)# glbp 1
switch(config-if-glbp)# weighting 110 lower 95 upper 105
switch(config-if-glbp)# weighting track 2 decrement 20
switch(config-if-glbp)# copy running-config startup-config
```

[Send document comments to nexus7k-docfeedback@cisco.com.](mailto:nexus7k-docfeedback@cisco.com)

Customizing GLBP

Customizing the behavior of GLBP is optional. Be aware that as soon as you enable a GLBP group by configuring a virtual IP address, that group is operational. If you enable a GLBP group before you customize GLBP, the router could take over control of the group and become the AVG before you finish customizing the feature. If you plan to customize GLBP, you should do so before enabling GLBP.

To customize GLBP, use the following commands in GLBP configuration mode:

| Command | Purpose |
|--------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>timers [msec] <i>hellotime</i> [msec] <i>holdtime</i></pre> <p>Example: switch(config-if-glbp)# timers 5 18</p> | <p>Configures the following hello and hold times for this GLBP member:</p> <ul style="list-style-type: none"> <i>hellotime</i>—The interval between successive hello packets sent by the AVG in a GLBP group. The range is from 1 to 60 seconds or from 250 to 60000 milliseconds. The default value is 3 seconds. <i>holdtime</i>—The interval before the virtual gateway and virtual forwarder information in the hello packet is considered invalid. The range is from 2 to 180 seconds or from 1020 to 180000 milliseconds. The default is 10 seconds. <p>The optional msec keyword specifies that the argument is expressed in milliseconds, instead of the default seconds.</p> |
| <pre>timers redirect <i>redirect</i> <i>timeout</i></pre> <p>Example: switch(config-if-glbp)# timers redirect 600 7200</p> | <p>Configures the following timers:</p> <ul style="list-style-type: none"> <i>redirect</i>—The time interval in seconds during which the AVG continues to redirect clients to an AVF. The range is from 0 to 3600 seconds. The default is 600 seconds. <i>timeout</i>—The interval in seconds before a secondary virtual forwarder becomes invalid. The range is from 610 to 64800 seconds. The default is 14,440 seconds. |
| <pre>priority <i>level</i></pre> <p>Example: switch(config-if-glbp)# priority 254</p> | <p>Sets the priority level used to select the AVG in a GLBP group. The range is from 1 to 255. The default is 100.</p> |
| <pre>preempt [delay minimum <i>seconds</i>]</pre> <p>Example: switch(config-if-glbp)# preempt delay minimum 60</p> | <p>Configures the router to take over as AVG for a GLBP group if it has a higher priority than the current AVG. This command is disabled by default.</p> <p>Use the optional delay minimum keywords and the <i>seconds</i> argument to specify a minimum delay interval in seconds before preemption of the AVG takes place.</p> <p>The seconds range is from 0 to 3600 seconds. The minimum delay default is 3600 seconds.</p> |

Send document comments to nexus7k-docfeedback@cisco.com.

Configuring Extended Hold Timers for GLBP

You can configure GLBP to use extended hold timers to support extended NSF during a controlled (graceful) switchover or ISSU, including software upgrades and supervisor switchovers. You should configure extended hold timers on all GLBP gateways (see the “[High Availability and Extended Nonstop Forwarding](#)” section on page 18-6).



Note

You must configure extended hold timers on all GLBP gateways if you configure non-default extended hold timers. You can configure different extended hold timer values on each GLBP gateway, based on the expected system switchover delays.

To configure GLBP extended hold timers, use the following command in global configuration mode:

| Command | Purpose |
|----------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|
| <code>glbp timers extended-hold [timer]</code> | Sets the GLBP extended hold timer, in seconds. The timer range is from 10 to 255. The default is 10. |
| Example: <pre>switch(config)# glbp timers extended-hold 30</pre> | |

Use the `show glbp` command to display the extended hold time.

Enabling a GLBP Group

You can configure the virtual IP address on an interface to enable the GLBP group. You must configure each gateway in the GLBP group with the same group number. The GLBP member can learn all other required parameters from another GLBP member.

BEFORE YOU BEGIN

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Enable GLBP (see the “[Enabling GLBP](#)” section on page 18-8).

SUMMARY STEPS

1. **configure terminal**
2. **interface** *interface-type slot/port*
3. **ip** *ip-address/length*
4. **glbp** *group-number*
5. **ip** [*ip-address* [**secondary**]]
6. (Optional) **show glbp** [*group group-number*] [**brief**]
7. (Optional) **copy running-config startup-config**

[Send document comments to nexus7k-docfeedback@cisco.com.](mailto:nexus7k-docfeedback@cisco.com)

DETAILED STEPS

| | Command | Purpose |
|--------|------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | configure terminal Example: switch# configure terminal switch(config)# | Enters configuration mode. |
| Step 2 | interface <i>interface-type slot/port</i> Example: switch(config)# interface ethernet 1/2 switch(config-if)# | Enters interface configuration mode. |
| Step 3 | ip <i>ip-address/length</i> Example: switch(config-if)# ip 192.0.2.1/8 | Configures the IPv4 address for the interface. |
| Step 4 | glbp <i>group-number</i> Example: switch(config-if)# glbp 1 switch(config-if-glbp)# | Creates a GLBP group and enters GLBP configuration mode. |
| Step 5 | ip [<i>ip-address</i> [secondary]] Example: switch(config-if-glbp)# ip 192.0.2.10 | Enables GLBP on an interface and identifies the virtual IP address. The virtual IP should be in the same subnet as the interface IP address. After you identify a virtual IP address, you can use the glbp group ip command again with the secondary keyword to indicate additional IP addresses supported by this group. If you only use the ip keyword, GLBP learns the virtual IP address from the neighbors. |
| Step 6 | show glbp [group <i>group-number</i>] [brief] Example: switch(config-if-glbp)# show glbp brief | (Optional) Displays a brief summary of GLBP information. |
| Step 7 | copy running-config startup-config Example: switch(config-if-glbp)# copy running-config startup-config | (Optional) Saves this configuration change. |

This example shows how to enable GLBP on Ethernet 1/2:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# glbp 1
switch(config-if-glbp)# ip 192.0.2.10
```


Send document comments to nexus7k-docfeedback@cisco.com.

Verifying the GLBP Configuration

To display GLBP configuration information, perform one of the following tasks:

| Command | Purpose |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| show glbp [group <i>group-number</i>] | Displays the GLBP status for all or one group. |
| show glbp capability | Displays the GLBP capability for all or one group. |
| show glbp interface <i>interface-type slot/port</i> | Displays the GLBP status for an interface. |
| show glbp interface <i>interface-type slot/port</i> [active] [disabled] [init] [listen] [standby] | Displays the GLBP status for a group or interface for virtual forwarders in the selected state. |
| show glbp interface <i>interface-type slot/port</i> [active] [disabled] [init] [listen] [standby] brief | Displays a brief summary of the GLBP status for a group or interface for virtual forwarders in the selected state. |

Configuration Examples for GLBP

The following example shows how to enable GLBP on an interface, with MD5 authentication, interface tracking, and weighted load balancing:

```
key chain glbp-keys
key 0
  key-string 7 zqdest
  accept-lifetime 00:00:00 Jun 01 2008 23:59:59 Sep 12 2008
  send-lifetime 00:00:00 Jun 01 2008 23:59:59 Aug 12 2008
key 1
  key-string 7 uaegdyito
  accept-lifetime 00:00:00 Aug 12 2008 23:59:59 Dec 12 2008
  send-lifetime 00:00:00 Sep 12 2008 23:59:59 Nov 12 2008
feature glbp
track 2 interface ethernet 2/2 ip
interface ethernet 1/2
  ip address 192.0.2.2/8
  glbp 1
    authentication md5 key-chain glbp-keys
    weighting 110 lower 95 upper 105
    weighting track 2 decrement 20
  ip 192.0.2.10
  no shutdown
```

Additional References

For additional information related to implementing GLBP, see the following sections:

- [Related Documents, page 18-18](#)
- [Standards, page 18-18](#)

[Send document comments to nexus7k-docfeedback@cisco.com.](mailto:nexus7k-docfeedback@cisco.com)

Related Documents

| Related Topic | Document Title |
|----------------------------------------------------|------------------------------------------------------------------------------------------|
| Configuring the Hot Standby Redundancy protocol | Chapter 19, “Configuring HSRP” |
| Configuring the Virtual Router Redundancy protocol | Chapter 20, “Configuring VRRP” |
| GLBP CLI commands | <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 5.x</i> |
| Configuring high availability | <i>Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide, Release 5.x</i> |

Standards

| Standards | Title |
|---------------------------------------------------------------------------------------------------------------------------------------|-------|
| No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature. | — |

Feature History for GLBP

[Table 18-2](#) lists the release history for this feature.

Table 18-2 Feature History for GLBP

| Feature Name | Releases | Feature Information |
|--------------|----------|------------------------------|
| GLBP | 4.0(1) | This feature was introduced. |