



network bgp commands

ONTAP 9.11.1 commands

NetApp

February 11, 2024

Table of Contents

- network bgp commands 1
 - network bgp config create 1
 - network bgp config delete 1
 - network bgp config modify 2
 - network bgp config show 2
 - network bgp defaults modify 3
 - network bgp defaults show 4
 - network bgp peer-group create 4
 - network bgp peer-group delete 6
 - network bgp peer-group modify 6
 - network bgp peer-group rename 7
 - network bgp peer-group show 8
 - network bgp vserver-status show 9

network bgp commands

network bgp config create

Create BGP configuration

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `network bgp config create` command is used to create the border gateway protocol (BGP) configuration for a node. It can be used to override the BGP parameters defined in the global BGP defaults.

Parameters

-node {<nodename>|local} - Node

This parameter specifies the node on which configuration details will be created.

-asn <integer> - Autonomous System Number

This parameter specifies the autonomous system number (ASN). The ASN attribute is a positive integer of the range from 1 to 4,294,967,295. It should typically be chosen from RFC6996 "Autonomous System (AS) Reservation for Private Use" or the AS number assigned to the operator's organization.

-hold-time <integer> - Hold Time

This parameter specifies the hold time in seconds. The default value is 180.

-router-id <IP Address> - Router ID

This parameter specifies the local router ID. The router-id value takes the form of an IPv4 address. The default router-id will be initialized using a local IPv4 address in admin vserver.

Examples

```
cluster1::> network bgp config create -node node1 -asn 10 -hold-time 180  
-router-id 10.0.1.112
```

network bgp config delete

Delete BGP configuration

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `network bgp config delete` command deletes a node's border gateway protocol (BGP) configuration. A BGP configuration cannot be deleted if there are BGP peer groups configured on the associated node.

Parameters

-node {<nodename>|local} - Node

This parameter specifies the node for which the BGP configuration will be deleted.

Examples

```
cluster1::> network bgp config delete -node node1
```

network bgp config modify

Modify BGP configuration

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `network bgp config modify` command is used to modify a node's border gateway protocol (BGP) configuration.

Parameters

-node {<nodename>|local} - Node

This parameter specifies the node on which BGP configuration will be modified.

[-asn <integer>] - Autonomous System Number

This parameter specifies the autonomous system number (ASN). The ASN attribute is a positive integer of the range from 1 to 4,294,967,295. It should typically be chosen from RFC6996 "Autonomous System (AS) Reservation for Private Use" or the AS number assigned to the operator's organization.

[-hold-time <integer>] - Hold Time

This parameter specifies the hold time in seconds.

[-router-id <IP Address>] - Router ID

This parameter specifies the local router ID. The router-id value takes the form of an IPv4 address.

Examples

```
cluster1::> network bgp config modify -node node1 -router-id 1.1.1.1 -asn  
20
```

network bgp config show

Display BGP configuration

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `network bgp config show` command displays the border gateway protocol (BGP) configuration for each node.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `-fields ?` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-node {<nodename>|local}] - Node

This parameter selects the BGP configurations that match the specified node.

[-asn <integer>] - Autonomous System Number

This parameter selects the BGP configurations that match the specified autonomous system number.

[-hold-time <integer>] - Hold Time

This parameter selects BGP configurations that match the specified hold time.

[-router-id <IP Address>] - Router ID

This parameter selects the BGP configurations that match the specified router ID.

Examples

```
cluster1::> network bgp config show
Autonomous
System      Hold Time
Node        Number      (seconds)  Router ID
-----
node1       10         180        10.0.1.112
```

network bgp defaults modify

Modify BGP defaults

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `network bgp defaults modify` command modifies the global defaults for border gateway protocol (BGP) configurations.

Parameters

[-asn <integer>] - Autonomous System Number

This parameter specifies the autonomous system number (ASN). The ASN attribute is a positive integer. It should typically be chosen from RFC6996 "Autonomous System (AS) Reservation for Private Use", or the AS number assigned to the operator's organization. The default ASN is 65501.

[-hold-time <integer>] - Hold Time

This parameter specifies the hold time in seconds. The default value is 180.

Examples

```
cluster1::> network bgp defaults modify -asn 20
```

network bgp defaults show

Display BGP defaults

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `network bgp defaults show` command displays the global defaults for border gateway protocol (BGP) configurations.

Examples

```
cluster1::> network bgp defaults show
Autonomous
System Number   Hold Time
      (Seconds)
-----
10              180
```

network bgp peer-group create

Create a new BGP peer group

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `network bgp peer-group create` command is used to create a border gateway protocol (BGP) peer group. A BGP peer group will advertise VIP routes for the list of vservers in the peer group's `vserver-list` using the BGP LIF of the peer group. A BGP peer group will advertise VIP routes to a peer router using the border gateway protocol. The address of the peer router is identified by the `peer-address` value.

Parameters

-ipspace <IPspace> - IPspace Name

This parameter specifies the IPspace of the peer group being created.

-peer-group <text> - Peer Group Name

This parameter specifies the name of the peer group being created.

-bgp-lif <lif-name> - BGP LIF

This parameter specifies the BGP interface (BGP LIF) of the peer group being created.

-peer-address <IP Address> - Peer Router Address

This parameter specifies the IP address of the peer router for the peer group being created.

[-peer-asn <integer>] - Peer Router Autonomous number

This parameter specifies the peer router autonomous system number (ASN) in the peer group being created. The default value is the value of the local node's ASN.

-route-preference <integer> - Route Preference

This parameter specifies the preference field in BGP update messages for VIP routes. If a router receives multiple VIP route announcements for the same VIP LIF from different BGP LIFs, it will install the one that has the highest preference value. The default route preference value is 100.

[-asn-prepend-type <ASN Prepend type>] - ASN prepend type (privilege: advanced)

This parameter specifies the ASN that will be prepended in the BGP attributes. The possible values are `local-asn` and `peer-asn`. The default behaviour is not to prepend any ASN.

[-asn-prepend-count <integer>] - ASN prepend count (privilege: advanced)

This parameter specifies the number of times ASN, as specified in `asn-prepend-type` will be prepended in the BGP path attributes. The default behaviour is not to prepend any ASN.

[-community <BGP community>,...] - BGP Community (privilege: advanced)

This parameter specifies the communities that will be included in the BGP path attributes. The default behaviour is not to include any community in BGP path attributes.

[-med <integer>] - Multi Exit Discriminator (privilege: advanced)

This parameter specifies the Multi Exit Discriminator (MED) attribute of BGP update messages, which can be used by routers for best path selection, in cases where more than one peer advertises the same route with similar attributes.

[-use-peer-as-next-hop {true|false}] - Use Peer Address As Next Hop

This parameter specifies whether the peer group uses the peer address as a next hop route. When the value is true, the peer address is used as the next hop router for packets sent from VIP LIFs via the port on which `bgp-lif` is configured. Internally, a default route with a gateway configured as the `peer-address` is added automatically on the node for all the Vservers in this peer group's IPspace. The route will be added for a Vserver only if it has a VIP LIF hosted on the current node of `bgp-lif`. Note that these automatically installed default routes are for VIP traffic; however, they can be used for non-VIP traffic as well if a Vserver hosts both VIP and non-VIP LIFs in the same subnet as `bgp-lif`. This route will have metric of 20 and will be used to forward traffic through the current port of `bgp-lif`. The default value of this parameter is false.

Examples

```
cluster1::> network bgp peer-group create -peer-group group1 -ipspace
Default -bgp-lif bgp_lif -peer-address 10.0.1.112
```

network bgp peer-group delete

Delete a BGP peer group

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `network bgp peer-group delete` command is used to delete border gateway protocol (BGP) peer group configuration.

Parameters

-ipspace <IPspace> - IPspace Name

This parameter specifies the IPspace of the BGP peer group being deleted.

-peer-group <text> - Peer Group Name

This parameter specifies the name of the BGP peer group being deleted.

Examples

```
cluster1::> network bgp peer-group delete -ipspace Default -peer-group
group1
```

network bgp peer-group modify

Modify a BGP peer group

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `network bgp peer-group modify` command is used to modify a border gateway protocol (BGP) peer group configuration.

Parameters

-ipspace <IPspace> - IPspace Name

This parameter specifies the IPspace of the peer group being modified.

-peer-group <text> - Peer Group Name

This parameter specifies the name of the peer group being modified.

[-peer-address <IP Address>] - Peer Router Address

This parameter specifies an updated value for the IP address of the peer router.

[-use-peer-as-next-hop {true|false}] - Use Peer Address As Next Hop

This parameter specifies whether the peer group uses the peer address as a next hop route. When the value is true, the peer address is used as the next hop router for packets sent from VIP LIFs via the port on which `bgp-lif` is configured. Internally, a default route with a gateway configured as the `peer-address` is added automatically on the node for all the Vservers in this peer group's IPspace. The route will be added for a Vserver only if it has a VIP LIF hosted on the current node of `bgp-lif`. Note that these automatically installed default routes are for VIP traffic; however, they can be used for non-VIP traffic as well if a Vserver hosts both VIP and non-VIP LIFs in the same subnet as `bgp-lif`. This route will have metric of 20 and will be used to forward traffic through the current port of `bgp-lif`. The default value of this parameter is false.

Examples

```
cluster1::> network bgp peer-group modify -ipspace Default -peer-group  
peer1 -peer-address 10.10.10.10
```

network bgp peer-group rename

Rename a BGP peer group

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `network bgp peer-group rename` command is used to assign a new name to a BGP peer group.

Parameters

-ipspace <IPspace> - IPspace Name (privilege: advanced)

This parameter specifies the IPspace of the peer group being renamed.

-peer-group <text> - Peer Group Name (privilege: advanced)

The name of the peer group to be updated.

-new-name <text> - New Name (privilege: advanced)

The new name for the peer group.

Examples

```
cluster1::> network bgp peer-group rename -peer-group old_name -new-name  
new_name
```

network bgp peer-group show

Display BGP peer groups information

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `network bgp peer-group show` command displays the BGP peer groups configuration.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-ipspace <IPspace>] - IPspace Name

This parameter selects peer groups that match the specified IPspace.

[-peer-group <text>] - Peer Group Name

This parameter selects peer groups that match the specified name.

[-bgp-lif <lif-name>] - BGP LIF

This parameter selects peer groups that match the specified BGP Interface.

[-peer-address <IP Address>] - Peer Router Address

This parameter selects peer groups that match the specified peer router address.

[-peer-asn <integer>] - Peer Router Autonomous number

This parameter selects peer groups that match the specified autonomous system number.

[-state <BGP Session State>] - Peer Group State

This parameter selects peer groups that match the specified BGP session state.

[-bgp-node <nodename>] - BGP LIF Node

This parameter selects peer groups that match the specified `bgp-node` value. This value is calculated based on the current node of the corresponding BGP LIF.

[-bgp-port <netport>] - BGP LIF Port

This parameter selects peer groups that match the specified `bgp-port` value. This value is calculated based on the current port of the associated BGP LIF.

[-route-preference <integer>] - Route Preference

This parameter selects peer groups that match the specified route preference value.

[`-asn-prepend-type` <ASN Prepend type>] - ASN prepend type (privilege: advanced)

This parameter selects peer groups that match the specified `asn-prepend-type` value. The possible values are `local-asn` and `peer-asn`.

[`-asn-prepend-count` <integer>] - ASN prepend count (privilege: advanced)

This parameter selects peer groups that match the specified `asn-prepend-count` value.

[`-community` <BGP community>,...] - BGP Community (privilege: advanced)

This parameter selects peer groups that match the specified `community` value.

[`-med` <integer>] - Multi Exit Discriminator (privilege: advanced)

This parameter selects peer groups that match the specified `med` value.

[`-use-peer-as-next-hop` {`true`|`false`}] - Use Peer Address As Next Hop

This parameter selects peer groups that match the specified `use-peer-as-next-hop` value.

Examples

```
cluster1::> network bgp peer-group show
IPspace: Default
Peer          Local BGP Peer router          Autonomous
Group         Interface Address/subnet    state          Number      Node
Port
-----
gp1           bgp_lif1  10.0.5.37         up              10
node1 e1a
gp2           bgp_lif2  10.0.6.38         up              12
node1 e2a
```

network bgp vserver-status show

Display Vserver BGP status

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `network bgp vserver-status show` command displays the per-node border gateway protocol (BGP) status for each vserver. The BGP status for a particular vserver is "up" when at least one BGP peer group supporting that vserver is able to communicate with its peer router.

Parameters

{ [`-fields` <fieldname>,...]

If you specify the `-fields` <fieldname>, ... parameter, the command output also includes the specified

field or fields. You can use '-fields ?' to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-node {<nodename>|local}] - Node

This parameter selects the BGP status that match the specified node.

[-vserver <vserver name>] - Vserver

This parameter selects the BGP status for specified vservers.

[-ipv4-status {unknown|unconfigured|up|down}] - IPv4 status

This parameter selects the BGP status that matches the specified status for IPv4 address family.

[-ipv6-status {unknown|unconfigured|up|down}] - IPv6 status

This parameter selects the BGP status that matches the specified status for IPv6 address family.

Examples

```
cluster1::> network bgp vservers-status show
Node          vservers  IPv4 status IPv6 status
-----
node1         vs1       up          up
```

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.