accepted-prefix-limit

Syntax

```
accepted-prefix-limit {
    maximum number;
    teardown <percentage-threshold> idle-timeout (forever | minutes);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols bgp family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems *logical-system-name* protocols bgp family route-target],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* family route-target],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address* family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address* family route-target],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp family route-target],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* family route-target],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address* family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address* family route-target],
[edit protocols bgp family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit protocols bgp family route-target],
[edit protocols bgp group *group-name* family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit protocols bgp group *group-name* family route-target],
[edit protocols bgp group *group-name* neighbor *address* family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit protocols bgp group *group-name* neighbor *address* family route-target],
[edit routing-instances *routing-instance-name* protocols bgp family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit routing-instances *routing-instance-name* protocols bgp family route-target],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* family route-target],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address* family (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address* family route-target]

Release Information Statement introduced in JUNOS Release 9.2.
Statement introduced in JUNOS Release 9.2 for EX Series switches.

Description Configure a limit to the number of prefixes that can be accepted on a BGP peering session. When that limit is exceeded, a system log message is sent. You can optionally specify to reset the BGP session when the number of accepted prefixes exceeds the specified limit.

Options idle-timeout (forever | *minutes*)—Specify that a BGP session that has been reset is not reestablished until after the specified timeout period. Specify forever to prevent the BGP session from being reestablished until the clear bgp neighbor command is issued.

maximum *number*—Limit the number of prefixes that can be accepted on a BGP peering session. A system log message is sent when that number is exceeded.

Range: 1 through 4,294,967,295 (2^{32} – 1)

teardown <*percentage* 1/n threshold>—Specify to reset the BGP peering session when the specified limit to the number of prefixes that can be accepted is exceeded. If you specify a percentage, a system log message is sent when the accepted number of prefixes on the BGP session exceeds the specified percentage of the configured limit. After a BGP session is reset, it is reestablished within a short time unless you include the idle-timeout statement.

Range: 1 through 100

Range: 1 through 2400

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Topics
- prefix-limit
- Enabling Multiprotocol BGP

Published: 2010-04-14