Uses of the Overload Bit with IS–IS

Document ID: 24509

Contents

Introduction

This document introduces the **set−overload−bit** Intermediate System−to−Intermediate System (IS−IS) configuration command, and how and when to use it with the **wait−for−bgp** and **suppress** keywords. Throughout this document, the term Intermediate System (IS) and router are interchangeable.

Prerequisites

Requirements

Readers of this document should have a basic knowledge of:

- Border Gateway Protocol (BGP) and IS−IS routing protocols.

Components Used

The information in this document is based on these software and hardware versions:

- Cisco IOS® Software Release 12.1(9)
- Cisco 2500 and 3600 Series routers

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Conventions

For more information on document conventions, refer to the Cisco Technical Tips Conventions.

Traditional Use of the Overload Bit

When a router runs out of system resources (memory or CPU), it cannot store the link−state database or run shortest path first (SPF). In this situation, the router should alert other routers within its area by setting a particular bit in its link−state packets (LSPs). When other routers detect that this bit is set, they will not use

this router for transit traffic but they will use it for packets destined to the overloaded router's directly connected networks and IP prefixes.

In IS–IS, a router immediately floods its own LSP even before sending complete sequence number PDU (CSNP) packets. The overload bit is thereby used to advise the rest of the network not to route transit traffic through the newly reloaded router.

For each LSP, the ISO/IEC 10589:1992 ⬈ defines a special bit called LSP Database Overload Bit. The draft mentions the overloaded condition (in section 7.3.19): "As a result of network mis–configuration, or certain transitory conditions, it is possible that there may be insufficient memory resources available to store a received Link State PDU. When this occurs, an IS needs to take certain steps to ensure that if its LSP database becomes inconsistent with the other ISs', that these ISs do not rely on forwarding paths through the overloaded IS."

When an IS is in this condition, it sets this bit in the non–pseudonode LSP fragment 0 that it generates.

Also, in the draft, section 7.2.8.1 makes a note that other ISs should not use the overloaded IS as a transit router, but may reach end systems (ESs) that are attached directly. During this time directly connected interfaces, as well as IP prefixes, are still reachable. Cisco IOS does not use the overload bit for this functionality, although the ability to permanently set the overload–bit was introduced into IOS with Cisco Bug ID CSCdj18100. In Cisco implementation, when overload bit set is set, the directly connected interfaces/IP prefixes are reachable.

Expanded Use of the Overload Bit

The IS–IS overload bit technique was expanded with Cisco bug ID CSCdp01872 (registered customers only) . You can configure a router to advertise its LSP with the overload bit for a specific amount of time after a reload. When the timer expires, the overload bit is cleared and the LSP is re–flooded.

This new functionality is useful to Internet Service Providers (ISPs) who run both Border Gateway Protocol (BGP) and IS–IS to avoid a couple of "black hole" scenarios. Setting the overload bit for a fixed amount of time right after a reload ensures that the router does not receive transit traffic while the routing protocol is still converging.

The technique for setting this bit for a certain period of time after a reload is implemented using the following command. This command takes a 5 to 86400 second range of time for the overload bit to remain set after reload.

```
router isis
set-overload-bit [on-startup [<timeout> | wait-for-bgp] ]
```

For example:

```
Router(config-router)#set-overload-bit on-startup 3500 wait-for-bgp

!--- Set the overload bit for 5 minutes (default is 10 minutes).
```

This feature also made it possible to configure a router to automatically disable the overload bit when BGP converged. For further information about waiting for BGP, please see RFC3277 Intermediate System to Intermediate System (IS–IS) Transient Blackhole Avoidance ⬈ .

According to the BGP specification, a BGP router does not need to send keepalives as long as it is sending updates. So, keepalives will be sent only after all updates have been sent. BGP is considered to have converged when keepalives are received from all the BGP neighbors.

If BGP keepalives are not received from all the BGP neighbors, and **wait−for−bgp** is configured, IS−IS will disable the overload bit after 10 minutes.

ISPs may want to suppress certain IP prefixes from being advertised in the router's own LSPs when **set−overload−bit** is configured. For example, it may not be desirable to allow Level 1 into Level 2 IP prefix propagation, which would make the router a transit node for IP traffic.

Cisco bug ID CSCdr98046 (registered customers only) gives more control over what happens when the the overload bit is being used in its expanded capacity. This enhancement allows an IS−IS Level 1 − Level 2 (L1L2) router which is redistributing IP routes from Level 1 into Level 2 or Level 2 into Level 1 to continue to advertise these redistributed routes in its LSP when the overload bit has been set.

Using the **suppress** keyword, you can configure a L1L2 router to redistribute and advertise IP routes from Level 1 into Level 2 or vice versa even when **set−overload−bit** is configured. The command syntax is as follows:

```
[no] set-overload-bit [on-startup [<n> | wait-for-bgp]] | [suppress [interlevel | external
```

The **suppress interlevel** keyword tells the router not to advertise IP prefixes learned from another IS−IS level if the overload bit is set. The **suppress external** keyword tells the router not to advertise IP prefixes learned from other protocols if the overload bit is set. The default is not to suppress and to maintain the Cisco bug ID CSCdp01872 (registered customers only) behavior.

The **suppress** option only takes effect when your own overload bit is set, and not when it is received or configured (for example, you could have set−overload−bit on−startup and the bit is not set).

```
router isis
set-overload-bit on-startup 40 suppress interlevel
```

In the above case, the overload bit is not actually set until the router is reloaded, and so you should continue to leak IP prefixes between levels. When you reload and actually set the bit, you should suppress the interlevel advertisements.

Configuration Example

The following network diagram is used to demonstrate the **set−overload−bit** command and the **wait−for−bgp** and **suppress** options.
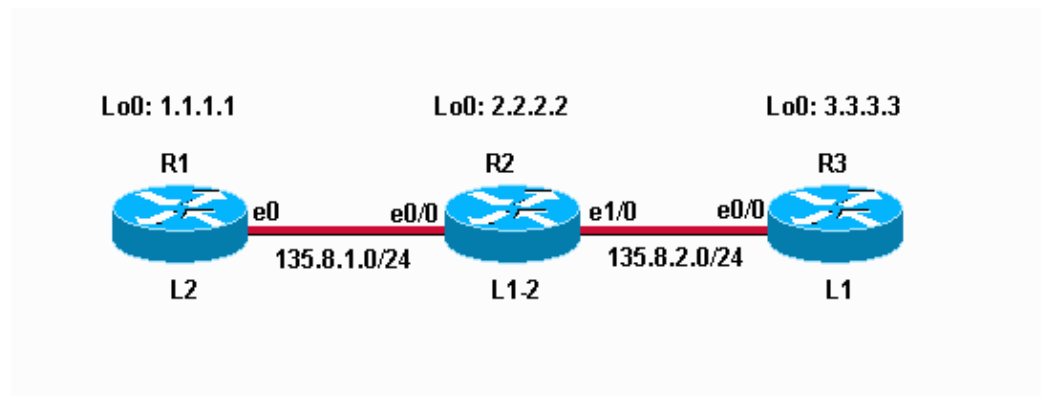
Here is the configuration containing the **wait−for−bgp** option on Router 2.

Router 2 Configuration

```
!
interface Loopback0
ip address 2.2.2.2 255.255.255.255

!--- Creates loopback interface and assigns
!--- IP address to interface Loopback0.

!
interface Ethernet0/0
ip address 135.8.1.1 255.255.255.0
ip router isis
!

!--- Assigns IP address to interface Ethernet0/0
!--- and enables IS-IS for IP on the interface.

!
!
interface Ethernet1/0
ip address 135.8.2.1 255.255.255.0
ip router isis
!

!--- Assigns IP address to interface Ethernet1/0
!--- and enables IS-IS for IP on the interface.

!
!
router isis
passive-interface Loopback0
net 12.0020.0200.2002.00
set-overload-bit on-startup wait-for-bgp
!

!--- Enables the IS-IS process on the router.
!--- Makes loopback interface passive
!--- (does not send IS-IS packets on interface).
!--- Assigns area and system ID to router.
!--- Sets the overload bit on startup to wait for BGP
!--- using the default timeout of 10 minutes.
```

The router has been freshly reloaded and before eBGP converged you can see that the overload bit is set on Router 2's LSP in the IS–IS Level 1 database.

```
    IS-IS Level-1 Link State Database:
    LSPID              LSP Seq Num  LSP Checksum  LSP Holdtime     ATT/P/OL
    r2.00-00           0x00000017   0x2372        284              0/0/1
```

Below, we see in the output of **debug isis update** that BGP has converged on Router 2:

```
    *Mar 1 00:00:51.015 UTC: BGP(0): Revise route installing 1.1.1.1/32
    -> 135.8.1.1 to main IP table
```

Router 2 now rebuilds its Level 1 LSP because BGP has converged and the overload bit is cleared. That is why you see "Important fields changed" in the output of **debug isis update** below.

```
    *Mar 1 00:00:51.087 UTC: ISIS-Upd: Building L1 LSP
    *Mar 1 00:00:51.087 UTC: ISIS-Upd: Important fields changed
    *Mar 1 00:00:51.087 UTC: ISIS-Upd: Full SPF required
```

Now we can see that Router 2 has completed its BGP update session with the neighbor:

```
     *Mar 1 00:00:52.127 UTC: BGP: 135.8.1.1 initial update completed
```

When we look at Router 2's Level 1 LSP again, we see that the Router 2 has cleared the overload bit (because BGP has converged) and that the LSP Seq Num field is increased by 1 (because a new LSP was created):

```
     IS-IS Level-1 Link State Database:
     LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime       ATT/P/OL
     r2.00-00             0x00000018   0xAD87        287                0/0/0
```

Here is the Router 2 configuration with L1L2 route leaking configured and the overload bit cleared.

Router 2 Configuration
| ```
!
interface Loopback0
ip address 2.2.2.2 255.255.255.255

!--- Creates loopback interface and assigns
!--- IP address to interface Loopback0.

!
interface Ethernet0/0
ip address 135.8.1.1 255.255.255.0
ip router isis

!--- Assigns IP address to interface Ethernet0/0
!--- and enables IS-IS for IP on the interface.

!
!
!
interface Ethernet1/0
ip address 135.8.2.1 255.255.255.0
ip router isis
!

!--- Assigns IP address to interface Ethernet1/0
!--- and enables IS-IS for IP on the interface.

!
!
router isis
redistribute static ip metric 11 level-1
redistribute isis ip level-2 into level-1 distribute-list 100
passive-interface Loopback0
net 12.0020.0200.2002.00
!

!--- Enables the IS-IS process on the router.
!--- Configured L2 to L1 route leaking
!--- Makes loopback interface passive
!--- (does not send IS-IS packets on interface).
!--- Assigns area and system ID to router.

!
ip route 200.200.200.200 255.255.255.255 loopback0

!--- Static route to 200.200.200.200 via loopback0.

access-list 100 permit ip any any

!--- Access list 100 is used to control which route
!--- gets leaked from Level 2 to Level 1.
``` |

Notice that Router 2's Level 1 database shows that the overload bit is clear in Router 2's L1 LSP.

```
IS-IS Level-1 LSP r2.00-00
LSPID        LSP Seq Num LSP Checksum LSP Holdtime ATT/P/OL
r2.00-00 * 0x0000005D   0xC252        180            0/0/0
Area Address: 12
NLPID: 0xCC
Hostname: r2
IP Address: 2.2.2.2
Metric: 10 IP 135.8.2.0 255.255.255.0
Metric: 10 IP 135.8.1.0 255.255.255.0
Metric:  0 IP 2.2.2.2 255.255.255.255
Metric: 10 IS r2.02
Metric: 10 IS r3.01
Metric: 11 IP-External 200.200.200.200 255.255.255.255
Metric:138 IP-Interarea 1.1.1.1 255.255.255.255
```

When we look at the IP routes that Router 3 is learning, we can see it learned Router 1's loopback address 1.1.1.1 from L2L1 route leaking. Notice also that Router 3 is also receiving the redistributed static route 200.200.200.0/32.

```
r3#show ip route isis
200.200.200.0/32 is subnetted, 1 subnets
i L1 200.200.200.200 [115/21] via 135.8.2.2, Ethernet0/0
1.0.0.0/32 is subnetted, 1 subnets
i ia 1.1.1.1 [115/148] via 135.8.2.2, Ethernet0/0
2.0.0.0/32 is subnetted, 1 subnets
i L1 2.2.2.2 [115/10] via 135.8.2.2, Ethernet0/0
135.8.0.0/24 is subnetted, 2 subnets
i L1 135.8.1.0 [115/20] via 135.8.2.2, Ethernet0/0
```

Now let's configure **set−overload−bit** on Router 2 with the **suppress** option. We will be suppressing both internal and external routes. The command syntax follows:

```
[no] set-overload-bit [on-startup [<n> | wait-for-bgp]] | [suppress [interlevel | external
```

suppress interlevel prevents the router from advertising prefixes learned from Level 2. **suppress external** prevents redistribution.

```
r2(config-router)#set-overload-bit suppress interlevel external
```

Looking into the Router 2 Level 1 database we can see the overload bit is now set in Router 2's Level 1 LSP. Both 200.200.200.200/32 and 1.1.1.1/32 have been suppressed. They are not injected into the Level 1 database.

```
IS-IS Level-1 LSP r2.00-00
LSPID        LSP Seq Num LSP Checksum LSP Holdtime ATT/P/OL
r2.00-00 * 0x0000005F   0x23C6        266            0/0/1
Area Address: 12
NLPID: 0xCC
Hostname: r2
IP Address: 2.2.2.2
Metric: 10 IP 135.8.2.0 255.255.255.0
Metric: 10 IP 135.8.1.0 255.255.255.0
Metric: 0 IP 2.2.2.2 255.255.255.255
Metric: 10 IS r2.02
Metric: 10 IS r3.01
```

When we enable **debug isis update−packets** on Router 2, we see "Important fields changed" in the output when the Level 1 and Level 2 LSPs are built. This indicates that the LSP content has changed, in other words, we have received an LSP that has the overload bit set. A new LSP requires a full SPF to be run.

```
*Mar 1 03:16:08.987 UTC: ISIS-Upd: Building L1 LSP
*Mar 1 03:16:08.987 UTC: ISIS-Upd: Important fields changed
*Mar 1 03:16:08.987 UTC: ISIS-Upd: Full SPF required
*Mar 1 03:16:08.987 UTC: ISIS-Upd: Building L2 LSP
*Mar 1 03:16:08.987 UTC: ISIS-Upd: Important fields changed
*Mar 1 03:16:08.987 UTC: ISIS-Upd: Full SPF required
*Mar 1 03:16:09.035 UTC: ISIS-Upd: Sending L1 LSP 0020.0200.2002.00-00, seq 61, ht 299 on
*Mar 1 03:16:09.095 UTC: ISIS-Upd: Sending L2 LSP 0020.0200.2002.00-00, seq 65, ht 299 on
```

Router 3's updated routing table no longer includes the IP networks 200.200.200.200 and 1.1.1.1.

```
r3#show ip route isis
2.0.0.0/32 is subnetted, 1 subnets
i L1 2.2.2.2 [115/10] via 135.8.2.2, Ethernet0/0
135.8.0.0/24 is subnetted, 2 subnets
i L1 135.8.1.0 [115/20] via 135.8.2.2, Ethernet0/0
```

DDTS Information

- Cisco bug ID CSCdj18100 (registered customers only) – Introduced the ability to set the overload bit manually.
- Cisco bug ID CSCdp01872 (registered customers only) – Introduced the ability to set the overload bit on startup. Wait until BGP has signaled convergence or set a timer to clear the overload bit.
- Cisco bug ID CSCdr98046 (registered customers only) – An IS−IS L1L2 router which is redistributing IP routes from Level 1 into Level 2 or Level 2 into Level 1 may continue to advertise these redistributed routes in its LSP when the overload bit has been set.

Related Information

- **IS−IS Support Page**
- **Routing Protocols Support Page**
- **Technical Support – Cisco Systems**

Updated: Oct 25, 2005   Document ID: 24509