



Border Gateway Protocol Link-State

Border Gateway Protocol Link-State (BGP-LS) is an Address Family Identifier (AFI) and Sub-address Family Identifier (SAFI) defined to carry interior gateway protocol (IGP) link-state database through BGP routing protocol. BGP-LS delivers network topology information to topology servers and Application Layer Traffic Optimization (ALTO) servers. BGP-LS allows policy-based control to aggregation, information-hiding, and abstraction. BGP-LS supports IS-IS and OSPFv2.

- [Information About Border Gateway Protocol Link-State, on page 1](#)
- [How to Configure OSPF With Border Gateway Protocol Link-State, on page 5](#)
- [How to Configure IS-IS With Border Gateway Protocol Link-State , on page 6](#)
- [Verifying Border Gateway Protocol Link-State Configurations, on page 7](#)
- [Border Gateway Protocol Link-State Debug Commands, on page 11](#)
- [Additional References for Border Gateway Protocol Link-State, on page 11](#)
- [Feature Information for Border Gateway Protocol Link-State, on page 12](#)

Information About Border Gateway Protocol Link-State

Overview of Link-State Information in Border Gateway Protocol

In a number of environments, a component external to a network is called upon to perform computations based on the network topology and current state of the connections within the network, including Traffic Engineering (TE) information. This information is typically distributed by interior gateway protocol (IGP) routing protocols within the network.

This module describes a mechanism by which link-state (LS) and Traffic Engineering (TE) information from IGPs are collected from networks and shared with external components using the BGP routing protocol, which uses a new BGP Network Layer Reachability Information (NLRI) encoding format. This mechanism is applicable to both physical and virtual links. Applications of this technique include Application-Layer Traffic Optimization (ALTO) servers and Path Computation Elements (PCEs), which are outside the network, but requires real-time information of the state of the network. For example, the link-state database information of each IGP node (OSPF or IS-IS) from the entire network.

In order to address the need for applications that require topological visibility across IGP areas, or even across Autonomous Systems (AS), the BGP-LS address-family or a sub-address-family have been defined to allow BGP to carry link-state information. The identifying key of each link-state object, for example, a node, link, or prefix, is encoded in the NLRI and the properties of the object are encoded in the BGP-LS attribute.

The below figure describes a typical deployment scenario of a network that utilizes BGP-LS. In each IGP area, one or more nodes are configured with BGP-LS. These BGP speakers form an IBGP mesh by connecting to one or more route-reflectors (RR). This way, all BGP speakers (specifically the route-reflectors (RR)) obtain link-state information from all IGP areas (and from other ASes from EBGp peers). An external component connects to the route-reflector to obtain this information (perhaps moderated by a policy regarding what information is or is not advertised to the external component). An external component (for example, a controller) then can collect these information in the "northbound" direction across IGP areas or ASes and construct the end-to-end path (with its associated SIDs) that are applied to an incoming packet for end-to-end forwarding.




Figure 1: Relation between IGP nodes and BGP


Carrying Link-State Information in Border Gateway Protocol

Carrying link-state information contains two parts:

- Definition of a new BGP NLRI that describes links, nodes, and prefixes comprises of IGP link-state information.
- Definition of a new BGP-LS attribute that carries link, node, and prefix properties and attributes, such as the link and prefix metric or auxiliary Router IDs of nodes, and so on.

TLV Format


Information in the new Link-State NLRIs and attributes is encoded in Type/Length/Value (TLV) triplets. The TLV format is shown in the below figure.



The Length field defines the length of the value portion in octets (thus, a TLV with no value portion would have a length of zero).

Link-State NLRI


The MP_REACH_NLRI and MP_UNREACH_NLRI attributes are BGP's containers for carrying opaque information. Each Link-State Network Layer Reachability Information (NLRI) describes either a node, a link, or a prefix. NLRI body is a set of Type/Length/Value triplets (TLV) and contains the data that identifies an object.



NLRI Types

The Total NLRI length field contains the cumulative length, in octets, of the rest of the NLRI, not including the NLRI Type field or itself.

Figure 2: The NLRI Types



The NLRI Types are shown in the following figures:

Figure 3: The Node NLRI Format





Figure 4: The Link NLRI Format



```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+
//                               Remote Node Descriptors (variable)      //
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
//                               Link Descriptors (variable)              //
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The IPv4 and IPv6 Prefix NLRIs (NLRI Type = 3 and Type = 4) use the same format, as shown in the following figure.

Figure 5: The IPv4/IPv6 Topology Prefix NLRI Format

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| Protocol-ID |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Identifier                               |
|                               (64 bits)                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
//                               Local Node Descriptors (variable)      //
+-----+-----+-----+-----+-----+-----+-----+-----+
//                               Prefix Descriptors (variable)           //
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Node Descriptors

Each link is anchored by a pair of Router-IDs that are used by the underlying IGP, namely, a 48-bit ISO System-ID for IS-IS and a 32-bit Router-ID for OSPFv2 and OSPFv3. An IGP may use one or more additional auxiliary Router-IDs, mainly for traffic engineering purposes. For example, IS-IS may have one or more IPv4 and IPv6 TE Router-IDs. These auxiliary Router-IDs must be included in the link attribute.

Link Descriptors

The Link Descriptor field is a set of Type/Length/Value (TLV) triplets. The link descriptor TLVs uniquely identify a link among multiple parallel links between a pair of anchor routers. A link described by the link descriptor TLVs actually is a "half-link", a unidirectional representation of a logical link. In order to fully describe a single logical link, two originating routers advertise a half-link each, that is, two Link NLRIs are advertised for a given point-to-point link.

Prefix Descriptors

The Prefix Descriptor field is a set of Type/Length/Value (TLV) triplets. Prefix Descriptor TLVs uniquely identify an IPv4 or IPv6 prefix originated by a node.

BGP-LS Attribute

The BGP-LS attribute is an optional, non-transitive BGP attribute that is used to carry link, node, and prefix parameters and attributes. It is defined as a set of Type/Length/Value (TLV) triplets. This attribute should only be included with Link-State NLRIs. This attribute must be ignored for all other address families.

How to Configure OSPF With Border Gateway Protocol Link-State

OSPF is one of the IGP protocols that feeds its topology into BGP into the LS cache. Link state information can be passed to BGP in two ways:

- When new communications between OSPF and BGP has been established, or when BGP-LS functionality has been initially enabled under OSPF, then all LSA information is downloaded to BGP via the LS library.
- As new LSA information is being processed or received from remote OSPF nodes, this information is added or updated in BGP.

Configuring Border Gateway Protocol Link-State With OSPF

Perform the following steps to configure OSPF with BGP-LS:

1. Enable the OSPF routing protocol and enter router configuration mode.

```
router ospf
```

For example,

```
Device(config-router)# router ospf 10
```

2. Distribute BGP link-state.

```
distribute link-state
```

For example,

```
Device(config-router)# distribute link-state instance-id <instid>
```

```
Device(config-router)# distribute link-state throttle <time>
```

instance-id (optional): Sets instance ID for LS distribution. Default Value is 0. Range: 32 to $2^{32}-1$.

throttle (optional): Sets throttle time to process LS distribution queue. Default value is 5 seconds. Range: 1 to 3600 seconds.



Note In the scenarios where any area gets deleted, throttle timer does not get honored. Queue is walked by OSPF completely and updates to all the areas are sent to BGP.

If you do not specify any value for instance ID and throttle, default values are taken.

Example:

```
#show run | sec router ospf
router ospf 10
distribute link-state instance-id 33 throttle 6
```



Note You should not be using the same instance ID for two OSPF instances. It throws an instance ID already in use error.

How to Configure IS-IS With Border Gateway Protocol Link-State

IS-IS distributes routing information into BGP. IS-IS processes the routing information in its LSP database and extract the relevant objects. It advertises IS-IS nodes, links, and prefix information and their attributes into BGP. This update from IS-IS into BGP only happens when there is a change in the LSP fragments, either belonging to the local router or any remote routers.

Configuring IS-IS With Border Gateway Protocol Link-State

Perform the following steps to configure IS-IS with BGP-LS:

1. Enable the IS-IS routing protocol and enter router configuration mode.

```
router isis
```

For example,

```
Device(config-router)# router isis
```

2. Distribute BGP link-state.

```
distribute link-state
```

For example,

```
Device(config-router)# distribute link-state instance-id <instid>
```

```
Device(config-router)# distribute link-state throttle <time>
```

instance-id (optional): Sets instance ID for LS distribution. The range is from 32-4294967294.

throttle (optional): Sets throttle time to process LS distribution queue. The range is from 5-20 seconds.

Configuring BGP

Perform the following steps to configure BGP with BGP-LS:

1. Enable the BGP routing protocol and enter router configuration mode.

```
router bgp
```

For example,

```
Device(config-if)# router bgp 100
```

2. Configure the address-family link-state.

```
address-family link-state link-state
```

For example,

```
Device(config-router)# address-family link-state link-state
```

3. Exit the address-family.

```
exit-address-family
```

For example,

```
Device(config-router)# exit-address-family
```

Example: Configuring ISIS With Border Gateway Protocol Link-State

Example: IS-IS Configuration

```
router isis 1
net 49.0001.1720.1600.1001.00
is-type level-1
metric-style wide
distribute link-state level-1
segment-routing mpls
segment-routing prefix-sid-map advertise-local
mpls traffic-eng router-id Loopback0
mpls traffic-eng level-1

interface GigabitEthernet2/2/2
ip address 172.16.0.1 255.255.0.0
ip router isis 1
negotiation auto
mpls traffic-eng tunnels
isis network point-to-point
```

Example: BGP Configuration

```
router bgp 100
bgp log-neighbor-changes
neighbor 10.0.0.1 remote-as 100
neighbor 10.0.0.4 remote-as 100
!
address-family ipv4
neighbor 10.0.0.1 activate
neighbor 10.0.0.4 activate
exit-address-family
!
address-family link-state link-state
neighbor 10.0.0.1 activate
neighbor 10.0.0.4 activate
exit-address-family
```

Verifying Border Gateway Protocol Link-State Configurations

Use the following show commands in any order to verify the status of the BGP-LS configurations.

show ip ospf ls-distribution

Displays the status of LS distribution.

```
Device# show ip ospf ls-distribution
```

```

      OSPF Router with ID (10.0.0.6) (Process ID 10)
      OSPF LS Distribution is Enabled
        Instance Id: 0
        Throttle time: 5
        Registration Handle: 0x0
        Status:Ready Active
      Num DBs Queued for LSCache Update: 0
      Num of DBs with Unresolved Links: 0

```

show ip ospf database dist-ls-pending

Displays the LSAs that are pending, to be sent to BGP.

Sample Output:

```
Device# show ip ospf database dist-ls-pending
```

```

      OSPF Router with ID (10.0.0.6) (Process ID 10)

      Router Link States (Area 0)

Link ID          ADV Router      Age          Seq#          Checksum Link count
10.0.0.7         10.0.0.6         4           0x80000006 0x009678 1
172.16.0.6       172.16.0.6       1110        0x80000018 0x00CAF9 2
(Has-unresolved-links)

```

show isis distribute-ls [level-1 | level-2]

Displays IS-IS internal LS cache information that are distributed to BGP.

```
Device# sh isis distribute-ls
```

```

ISIS distribute link-state: configured
distls_levels:0x3, distls_initialized:1,
distls_instance_id:0, distls_throttle_delay:10
LS DB: ls_init_started(0) ls_initialized(1) ls_pending_delete(0)
distls_enabled[1]:1
distls_enabled[2]:1
Level 1:
Node System ID:0003.0003.0003 Pseudonode-Id:0 ls_change_flags:0x0
LSP: lspid(0003.0003.0003.00-00), lsptype(0) lsp_change_flags(0x0)
Node Attr: name(r3) bitfield(0xD1) node_flags(0x0)
area_len/area_addr(2/33) num_mtid/mtid(0/0) ipv4_id(172.16.0.9)
num_alg/sr_alg(0/0) num_srgb/srgb(1/(start:16000, range:8000))
srgb_flags(0x80)
opaque_len/opaque(0/0x0)
ISIS LS Links:
mtid(0): nid:0002.0002.0002.00, {0, 0}, {6.6.6.1, 6.6.6.6}
Link Attr: bitbfield:0x940F, local_ipv4_id:6.6.6.1, remote_ipv4_id:172.16.0.8,
max_link_bw:10000, max_resv_bw:10000,
num_unresv_bw/unresv_bw:8/
[0]: 10000 kbits/sec, [1]: 8000 kbits/sec
[2]: 8000 kbits/sec, [3]: 8000 kbits/sec
[4]: 8000 kbits/sec, [5]: 8000 kbits/sec
[6]: 8000 kbits/sec, [7]: 8000 kbits/sec,
admin_group:0, protect_type:0, mpls_proto_mask:0x0,
te_metric:0, metric:0, link_name:,
num_srlg/srlg:0/
num_adj_sid/adjsid:2/
Adjacency SID Label:16 F:0 B:0 V:1 L:1 S:0 weight:0
Adjacency SID Label:17 F:0 B:1 V:1 L:1 S:0 weight:0
opaque_len/opaque_data:0/0x0
Address-family ipv4 ISIS LS Prefix:

```



```

mtid(0): 1.1.1.0/24
Prefix Attr: bitfield:0x0, metric:10, igp_flags:0x0,
  num_route_tag:0, route_tag:0
  num_pfx_sid:0, pfx_sid:
  pfx_srms:
  opaque_len:0, opaque_data:0x0
mtid(0): 172.16.0.8/24
Prefix Attr: bitfield:0x0, metric:10, igp_flags:0x0,
  num_route_tag:0, route_tag:0
  num_pfx_sid:0, pfx_sid:
  pfx_srms:
  opaque_len:0, opaque_data:0x0

```

show bgp link-state link-state

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
 r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
 x best-external, a additional-path, c RIB-compressed,
 t secondary path,

Origin codes: i - IGP, e - EGP, ? - incomplete

RPKI validation codes: V valid, I invalid, N Not found

Prefix codes: E link, V node, T4 IPv4 reachable route, T6 IPv6 reachable route, I Identifier,

N local node, R remote node, L link, P prefix,
 L1/L2 ISIS level-1/level-2, O OSPF, a area-ID, l link-ID,
 t topology-ID, s ISO-ID, c confed-ID/ASN, b bgp-identifier,
 r router-ID, i if-address, n nbr-address, o OSPF Route-type,
 p IP-prefix, d designated router address, u/U Unknown,
 x/X Unexpected, m/M Malformed

Network	Next Hop	Metric	LocPrf	Weight	Path
*> [V] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.1001.00]]	15.0.0.1	0		0 100	i
*> [V] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.2002.00]]	15.0.0.1	0		0 100	i
*> [V] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.3003.00]]	15.0.0.1	0		0 100	i
*> [V] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.4004.00]]	15.0.0.1	0		0 100	i
*> [V] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.5005.00]]	15.0.0.1	0		0 100	i
*>					
[E] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.1001.00]] [R[c100] [b0.0.0.0] [s1720.1600.2002.00]] [L]	15.0.0.1	0		0 100	i
*>					
[E] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.2002.00]] [R[c100] [b0.0.0.0] [s1720.1600.1001.00]] [L]	15.0.0.1	0		0 100	i
*>					
[E] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.2002.00]] [R[c100] [b0.0.0.0] [s1720.1600.3003.00]] [L]	15.0.0.1	0		0 100	i
*>					
[E] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.2002.00]] [R[c100] [b0.0.0.0] [s1720.1600.4004.00]] [L]	15.0.0.1	0		0 100	i
*>					
[E] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.3003.00]] [R[c100] [b0.0.0.0] [s1720.1600.2002.00]] [L]	15.0.0.1	0		0 100	i
*>					
[E] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.3003.00]] [R[c100] [b0.0.0.0] [s1720.1600.5005.00]] [L]					

```

15.0.0.1          0          0 100 i
*>
[E] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.4004.00]] [R[c100] [b0.0.0.0] [s1720.1600.2002.00]] [L]

15.0.0.1          0          0 100 i
*>
[E] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.4004.00]] [R[c100] [b0.0.0.0] [s1720.1600.5005.00]] [L]

15.0.0.1          0          0 100 i
*>
[E] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.5005.00]] [R[c100] [b0.0.0.0] [s1720.1600.3003.00]] [L]

15.0.0.1          0          0 100 i
*>
[E] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.5005.00]] [R[c100] [b0.0.0.0] [s1720.1600.4004.00]] [L]

15.0.0.1          0          0 100 i
*> [T4] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.1001.00]] [P[p10.0.0.0/24]]
15.0.0.1          0          0 100 i
*> [T4] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.1001.00]] [P[p7.7.7.7/32]]
15.0.0.1          0          0 100 i
*> [T4] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.2002.00]] [P[p10.0.0.0/24]]
15.0.0.1          0          0 100 i
*> [T4] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.2002.00]] [P[p11.0.0.0/24]]
15.0.0.1          0          0 100 i
*> [T4] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.2002.00]] [P[p12.0.0.0/24]]
15.0.0.1          0          0 100 i
*> [T4] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.2002.00]] [P[p5.5.5.5/32]]
15.0.0.1          0          0 100 i
*> [T4] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.3003.00]] [P[p11.0.0.0/24]]
15.0.0.1          0          0 100 i
*> [T4] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.3003.00]] [P[p13.0.0.0/24]]
15.0.0.1          0          0 100 i
*> [T4] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.3003.00]] [P[p3.3.3.3/32]]
15.0.0.1          0          0 100 i
*> [T4] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.4004.00]] [P[p12.0.0.0/24]]
15.0.0.1          0          0 100 i
*> [T4] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.4004.00]] [P[p14.0.0.0/24]]
15.0.0.1          0          0 100 i
*> [T4] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.4004.00]] [P[p15.15.15.15/32]]
15.0.0.1          0          0 100 i
*> [T4] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.5005.00]] [P[p13.0.0.0/24]]
15.0.0.1          0          0 100 i
*> [T4] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.5005.00]] [P[p14.0.0.0/24]]
15.0.0.1          0          0 100 i
*> [T4] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.5005.00]] [P[p15.0.0.0/24]]
15.0.0.1          0          0 100 i
*> [T4] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.5005.00]] [P[p16.16.16.16/32]]
15.0.0.1          0          0 100 i

```

show bgp link-state link-state nlri <nlri string>

```

BGP routing table entry for [V] [L1] [I0x43] [N[c100] [b0.0.0.0] [s1720.1600.4004.00]], version
95
Paths: (1 available, best #1, table link-state link-state)
  Not advertised to any peer
  Refresh Epoch 4
  Local
    16.16.16.16 (metric 30) from 15.15.15.15 (15.15.15.15)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Originator: 16.16.16.16, Cluster list: 15.15.15.15
      LS Attribute: Node-name: R4, ISIS area: 49.12.34
      rx pathid: 0, tx pathid: 0x0

```

Border Gateway Protocol Link-State Debug Commands

- **debug ip ospf dist-ls [detail]**

Turns on ls-distribution related debugs in OSPF.

- **debug isis distribute-ls**

Displays the items being advertised into the BGP from IS-IS.

Additional References for Border Gateway Protocol Link-State

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Commands List, All Releases

MIBs

MIB	MIBs Link
<ul style="list-style-type: none">• OSPF	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFCs	Title
<ul style="list-style-type: none">• RFC 7752	<i>Link-State Info Distribution Using BGP</i>

Technical Assistance

Description	Link
<p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p>	http://www.cisco.com/cisco/web/support/index.html

Feature Information for Border Gateway Protocol Link-State

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for BGP-LS

Feature Name	Releases	Feature Information
Border Gateway Protocol Link-State	Cisco IOS XE Everest 16.4.1	<p>BGP Link-State (LS) is an Address Family Identifier (AFI) and Sub-address Family Identifier (SAFI) defined to carry interior gateway protocol (IGP) link-state database through BGP. The following commands were introduced or modified:</p> <p>address-family link-state link-state, distribute link-state, show bgp link-state link-state, show bgp link-state link-state nlri <i>nlri string</i>, show ip ospf database dist-ls-pending, show ip ospf ls-distribution, show isis distribute-ls</p>