BGP Link-State Mechanisms and Update Groups

This chapter covers BGP Link-State fundamentals, configuration tasks for exchanging and distributing IGP link-state information, and update group mechanisms for optimizing BGP update generation and neighbor management.

BGP Link-State

BGP Link-State (BGP-LS) is a protocol extension that

- distributes IGP link-state information via BGP
- improves network topology visibility for applications like SR-PCE, and
- uses a dedicated AFI/SAFI (RFC7752) to encode link-state attributes.

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
BGP Link-State	Release 24.4.1	Introduced in this release on: Fixed Systems(8700)(select variants only*) With BGP Link-State (BGP-LS), you can efficiently share IGP link-state information across your network, allowing applications such as Segment Routing Path Computation Element (SR-PCE) to gain greater topology awareness and optimize path computations using Segment Routing Traffic Engineering (SR-TE). This feature uses Address Family Identifier (AFI) and Sub-address Family Identifier (SAFI) to encode link-state data in the BGP-LS attribute as defined by RFC7752. * BGP-LS functionality is now supported on Cisco-8712-MOD-M routers.

BGP-LS extends BGP to carry Interior Gateway Protocol (IGP) link-state information using a specific Address Family Identifier (AFI) and Sub-address Family Identifier (SAFI). As defined in RFC7752, BGP-LS encodes

each link-state object, such as a node, a link, or a prefix, in the BGP Network Layer Reachability Information (NLRI) format, while the properties of each object are conveyed using the BGP-LS path attribute. This approach enables controllers and applications to build a comprehensive topology view across multiple domains.

Example

A Segment Routing Path Computation Element (SR-PCE) can use BGP-LS data to discover node capabilities, learn mappings for SR segments, and compute optimal paths using Segment Routing Traffic Engineering (SR-TE). This enables SR-PCE to steer traffic on paths different from the underlying IGP-based distributed best-path computation.

Deployment scenarios for BGP Link-State

This figure illustrates a typical BGP-LS deployment.

Figure 1: Deployment scenario for BGP-LS

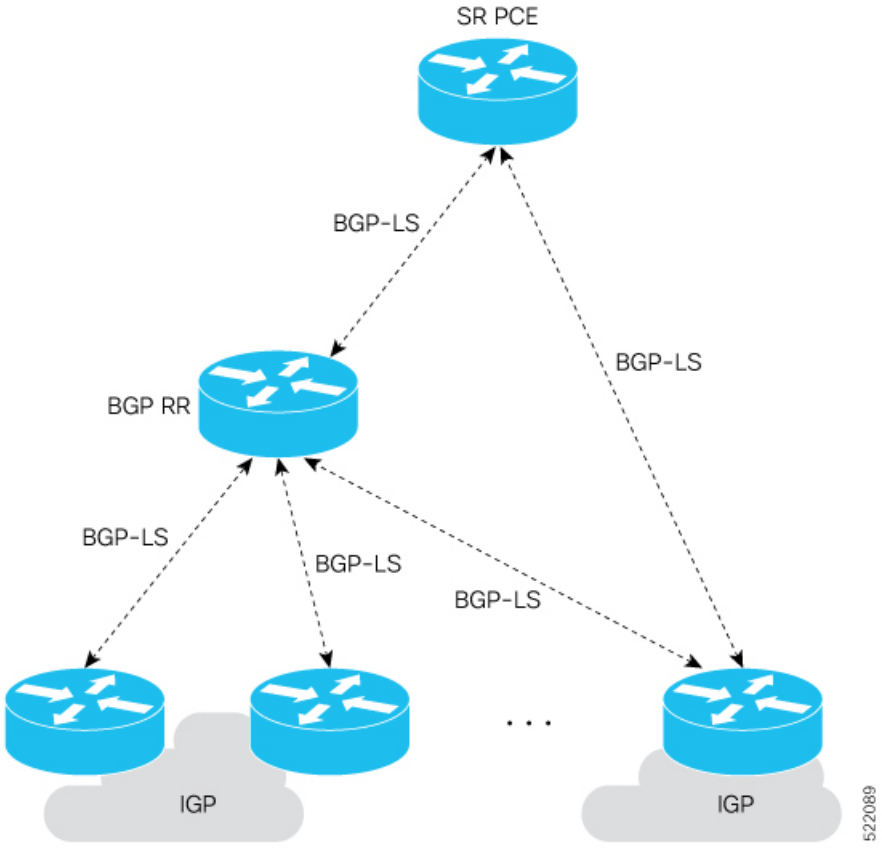

The deployment involves:

- Configuring one or more BGP speakers (nodes) with BGP-LS in each IGP area.
- Establishing an iBGP mesh between BGP speakers and route-reflectors.
- Allowing route-reflectors to aggregate and share link-state information from all IGP areas and from eBGP peers in other autonomous systems (AS).

Usage guidelines and limitations for BGP Link-State

Functional limitations

IGPs do not use BGP-LS data from remote BGP peers, and BGP does not download the received BGP-LS information into other router components.

Instance-id usage

The identifier field in BGP Link-State, called the instance-id, specifies the IGP routing domain to which the NLRI is associated.

Some best practices for instance-id configuration are:

- Assign a consistent instance-id value to all BGP-LS producers within a single IGP domain.
- When only one protocol instance is present, configure the instance-id value to 0.
- Use unique instance-id values for different routing protocol instances operating in separate IGP domains.

NLRIs with different instance-id values represent different IGP routing instances. If consistent instance-ids are not used, BGP-LS consumers may see duplicate objects or incorrect topology.

Configure BGP Link-State with a neighbor

Enable the exchange of BGP Link-State information with a BGP neighbor.

Use this task to distribute IGP link-state data (OSPF or IS-IS) to a BGP neighbor for use by controllers or applications such as SR-PCE.

Before you begin

- Ensure you have access to the router CLI.
- Confirm that the neighbor uses a private IP address.

Procedure

Step 1 Run the **configure** comamnd to enter configuration mode.

Example:

```
Router# configure
```

Step 2 Specify the BGP AS number and enter BGP configuration mode.

Example:

```
Router(config)# router bgp 100
```

Step 3 Configure the neighbor using its IP address.

Example:

```
Router(config-bgp)# neighbor 10.0.0.2
```

Step 4 Set the remote AS number for the neighbor and enter the link-state address family.

Example:

```
Router(config-bgp-nbr)# remote-as 1
Router(config-bgp-nbr)# address-family link-state link-state
Router(config-bgp-nbr)# commit
```

The router exchanges BGP Link-State data with the specified neighbor.

Configure a unique domain distinguisher (four-octet ASN)

Assign a unique identifier (domain distinguisher) for BGP Link-State using a four-octet ASN.

The domain distinguisher helps differentiate routing domains when distributing link-state information.

Follow these steps to configure the domain distinguisher.

Procedure

Step 1 In configuration mode, specify the BGP AS number and enter BGP configuration mode.

Example:

```
Router# configure
Router(config)# router bgp 100
```

Step 2 Enter the link-state address family configuration mode.

Example:

```
Router(config-bgp)# address-family link-state link-state
```

Step 3 Assign a unique domain distinguisher.

Example:

```
Router(config-bgp-af)# domain-distinguisher 1234
Router(config-bgp-af)# commit
```

Possible range of domain-distinguisher: 1 to 4294967295

The router uses the specified four-octet ASN as the domain distinguisher for BGP-LS.

Distributing IGP link-state databases with BGP-LS

BGP Link-State (BGP-LS) enables the distribution of IGP link-state databases, including OSPF and IS-IS link-state data, across multiple, independent routing domains. This distribution allows controllers or applications to build end-to-end paths that span several domains.

To distribute OSPFv2 link-state data using BGP-LS:

- Enter router configuration mode.
- Configure the OSPF process.

- Use the **distribute link-state** command with the appropriate instance-id.

Example configuration:

```
Router# configure
Router(config)# router ospf 100
Router(config-ospf)# distribute link-state instance-id 32
```

BGP update groups

A BGP update group is a mechanism that

- groups BGP neighbors with similar outbound routing policies
- enables efficient update message generation, and
- automatically recalculates group membership when configuration changes occur.

BGP update group generation separates update message creation from neighbor configuration by using an algorithm to assign neighbors to groups based on their outbound policies. This process operates automatically without requiring manual configuration. With update group-based message generation, the router sends a single update message to all neighbors in a group, which simplifies management, reduces redundant updates, and improves overall network performance.

How BGP update group membership and optimization work

Summary

The key components involved in the process are:

- Router: Applies outbound policies and manages update groups.
- BGP neighbors: Are assigned to update groups based on their outbound routing policies.
- Update group algorithm: Dynamically calculates group membership and manages update message generation.

BGP update group membership is recalculated automatically whenever there is a configuration change. The router ensures that neighbors with matching outbound routing policies are grouped together, which helps reduce processing and optimize updates.

Workflow

These stages describe how BGP update group membership and optimization work.

- The router detects a configuration change.
- The update group algorithm recalculates group memberships based on current outbound policies.
- The router generates update messages for each group and sends them to the appropriate neighbors.

Result

The router delivers optimized BGP updates, reducing redundancy and improving overall network efficiency.