CHAPTER 2

Overview of the Service Control Solution for MPLS/VPN Networks

Service Control in the MPLS/VPN Environment

MPLS/VPN networks are very complex and utilize many routing protocols and many different levels of addressing and control. In addition, the various VPNs may use overlapping IP addresses (private IPs).

The SCE platform makes a distinction between identical IP addresses that come from different VPNs, and maps them into subscribers according to the MPLS labels attached to the packets. This involves various mechanisms in all levels of the system.

The following assumptions and requirements allow the SCE platform to operate in an MPLS/VPN environment:

- The MPLS/VPN architecture is according to RFC-2547.
- The specific type of encapsulation used is the MPLS shim header over Ethernet (described in RFC-3032).
- There are two levels of MPLS labels.
  - External labels — Used for transport over the service provider MPLS core network.

  These labels are not mandatory for VPN classification, and some situations do not appear in the packet due to PHP or other reasons.
  - Internal labels (BGP labels) — Used to identify the VPNs connected to each edge router, and typically controlled by the BGP protocol.

  These labels are mandatory for VPN classification.
- The MPLS/VPN solution contains the SCE platform and the SM. The SM acts as a BGP peer for the PE routers in the service provider network, and communicates the BGP information to the SCE platform as subscriber information.

Note The MPLS/VPN solution supports the existence of non-VPN-based subscribers concurrently with the MPLS/VPN-based subscribers (see Non-VPN-Based Subscribers, page 2-6).

Definitions and Acronyms

The following table defines important terms and acronyms.

Table 2-1 MPLS/VPN Terms and Acronyms

Term or Acronym	Definition
PE (Provider Edge router)	A router at the edge of the service provider network. The PE routers are the ones that connect to the customers, and maintain the VPNs
P (Provider router)	A router in the core of the service provider network. P routers only forward MPLS packets, regardless of VPNs.
VPN (Virtual Private Network)	In the Service Control context, a VPN is the part of the VPN that resides in a specific site. It is a managed entity over which private IP subscribers can be managed.
BGP LEG	A software module that resides on the SM server and generates BGP-related login events. The BGP LEG communicates with the BGP routers (PEs) and passes the relevant updates to the SM software, which generates login events to the SCE platform for the updated VPN-based subscribers.
Upstream	Traffic coming from the PE router and going into the P router
Downstream	Traffic coming from the P router and going into the PE router
RD (Route Distinguisher)	Used to uniquely identify the same network/mask from different VRFs (such as, 10.0.0.0/8 from VPN A and 10.0.0.0/8 from VPN B)
RT (Route Target)	Used by the routing protocols to control import and export policies, to build arbitrary VPN topologies for customers
VRF (Virtual Routing and Forwarding instance)	Mechanism used to build per-interface routing tables. Each PE has several VRFs, one for each site it connects to. This is how the private IPs remain unique.

What are the Challenges for Service Control for MPLS/VPN Support?

- Private IP addresses cause flows to look the same except for their MPLS labels.
- The MPLS labels are different in each direction, and must be matched.
- Detecting that a flow belongs to a certain VPN is complicated by the fact that in the downstream direction there is no external label. The SCE platform must be able to understand the VPN information from the internal label + the MAC address of the PE..

How MPLS/VPN Support Works

Service Control supports three mechanisms that make MPLS/VPN support work:

- Flow detection – This is the job of the SCE platform, to match upstream and downstream traffic to identify flows.
- VPN detection – Downstream VPN labels are identified by the SM. The SCE platform learns the upstream labels from the traffic to identify the VPN.
- Subscriber detection – The SM and the SCE platform function together to identify the IP range within a VPN that is defined as a single subscriber.

Flow Detection

Flow detection is the process of deciding which packets belong to the same flow. This relates to the first two challenges listed:

- Private IP addresses cause flows to look the same except for their MPLS labels.
- The MPLS labels are different in each direction, and must be matched.

Flow detection is based on the MPLS labels, extending the basic 5 tuple that SCOS uses to identify flows, and taking into account the fact that in MPLS, the packet is labeled differently in each direction.

Since MPLS traffic is unidirectional, each direction is classified separately by the SCE platform, using the following:

- Downstream – the BGP label and the MAC address of the PE (only one label that is relevant to the classification)

  Downstream labels are learned from the control plane (through the SM BGP LEG).
- Upstream – the combination of the external label, the BGP label, and the MAC address of the P router (two labels that are relevant to the classification)

  Upstream labels are learned from the data plane.

VPN Detection

The network configuration that provides the division into VPNs is controlled by the SM. The network-wide value that describes a VPN most closely is either the Route Target or the Route Distinguisher

- The administrator configures the SM to detect VPNs, according to selected attribute (RT or RD) (see How to Configure the SM for MPLS/VPN Support, page 3-6.)
- The network operator provides the SCE platform with a mapping between RT values and VPN subscriber names. (See How to Manage MPLS/VPN Support via SM CLU, page 4-8)

The relevant module in the Subscriber Manager server (SM) is the BGP-LEG. The BGP-LEG is added to the BGP neighborhood for obtaining the information on the MPLS labels. The local PEs are configured to add the BGP-LEG as a BGP peer.

The SCE platform detects that a flow belongs to a certain VPN according to the downstream label that the flow carries, and the MAC address of the PE that it is sent to.

One VPN may spread over more than one PE router, as long as all the sites of the VPN are connected to the subscriber side of the same SCE platform

VPNs can be configured only via the SM. The SCE platform CLI can be used to view VPN-related information, but not to configure the VPNs.

Subscriber Detection

What is an MPLS/VPN-based Subscriber?

As in other modes of operation, in MPLS/VPN each flow belongs to a certain subscriber. A VPN-based subscriber is a part of a VPN. The VPN itself corresponds to a set of IP addresses that are managed separately and that belong to a specific ISP customer who pays for the VPN service.

An MPLS/VPN-based subscriber can be defined as either of the following:

- A set of IP addresses or ranges in a certain VPN.
- All the IP addresses of a CE router, defined by a BGP community over a VPN.

The network configuration that provides the division into VPNs and VPN-based subscribers is controlled by the SM. (For more information, see the *Cisco Service Control Management Suite Subscriber Manager User Guide*)

Private IP Subscriber Support

VPN-based subscribers can have private IP mappings, which are a combination of an IP range and a VPN mapping. Since the source of such mappings is typically in the BGP protocol, and they are received automatically from the protocol by the BGP agent, the IP ranges may contain overlapping ranges. The semantics of such overlaps is that of a longest prefix match.

For example, if subscriber A receives the range 10.0.0.0/8@VPN1 and subscriber B receives the range 10.1.0.0/16@VPN1, then the system maps IPs that start with 10.1 to subscriber B, and any other address that begins with 10 to subscriber A. Traffic with other IP addresses on VPN1 will be mapped to the unknown subscriber.

For private IP subscribers, flows are distributed to traffic processors according to the VPN, not according to the IP address. This means that all traffic from any one VPN is mapped to the same traffic processor.

How the Service Control MPLS/VPN Solution Works

How the Service Control MPLS/VPN Solution Works: A Summary

- The SM is configured with the VPNs and VPN-based subscribers that should be managed.

  A VPN is identified by the RD / RT and the PE.
- The BGP-LEG updates the SM with the MPLS labels and IP routes.
- The SM pushes the VPNs with their labels and the VPN-based subscriber to the SCE platform with the downstream MPLS labels of the VPN.
- The SCE platform resolves the PE MAC addresses and updates its tables with the new information.
- The SCE platform learns the upstream labels, including the P MAC address.
- The SCE platform provides the regular services to the VPN-based subscribers (BW management, reports, etc.)

SCE Platform Tasks in the MPLS/VPN Solution

- Matching upstream to downstream labels
  - Mappings of downstream labels to VPNs are received from the SM
  - Upstream labels are learned from the data
- The MAC addresses of the PEs are used to distinguish downstream labels of different PEs
- After the learning, each flow is classified as belonging to one of the VPNs.
- The SCE platform performs a longest prefix match on the IP address inside the VPN, and classifies each flow to the correct VPN-based subscriber
- The SCE platform runs the SCA-BB application for the network flows, which are classified to VPNs, thus providing subscriber aware service control and reporting

BGP LEG Tasks in the MPLS/VPN Solution

- The BGP LEG is a software module that runs on the SM server
- The LEG maintains a BGP session with a list of PEs
- After the sessions establishment, the LEG propagates MP-BGP route-updates from the PEs to the SM module

SM Tasks in the MPLS/VPN Solution

- The VPNs are stored in the SM database.
- Each VPN is defined by:

  - The IP address of the loopback interface of the PE router.
  - The RD or RT that identifies the VPN within the PE router.
- A VPN-based subscriber is defined by the IP range in a specified VPN or the BGP community (CE as subscriber).
- The SM receives updates from the BGP LEG, and updates the VPN information with the new MPLS labels.
- The relevant SCE platforms that will get the MPLS updates are defined by the VPN domain.

Service Control MPLS/VPN Concepts

Non-VPN-Based Subscribers

The MPLS/VPN solution supports the existence of non-VPN-based (regular IP) subscribers concurrently with the MPLS/VPN-based subscribers, with the following limitations and requirements:

- The SM must work in "push" mode.
- Non-VPN-based subscribers cannot have IP in VPN mappings.
- VLAN-based subscribers are NOT supported at the same time as MPLS/VPN-based subscribers.

In typical MPLS/VPN networks, traffic that does not belong to any VPN is labeled with a single MPLS label in the upstream direction, which is used for routing. The downstream direction of such flows typically contains no label, due to penultimate hop popping.

The SCE platform uses the one or more labels upstream and no label downstream definition to identify non-VPN flows. Classification and traffic processor load balancing on these flows is performed according to the IP header, rather than the label.

This process requires learning of the upstream labels in use for such flows, and is done using the flow detection mechanism described above (see Flow Detection, page 2-3).

Bypassing Unknown VPNs

In an MPLS network, there may be many VPNs crossing the SCE platform, only a small number of which require service control functionality. It is necessary for the SCE platform to recognize which VPNs are not managed.

- The SCE platform automatically bypasses any VPN that is not configured in the SM
- The VPNs are bypassed by the SCE platform without any service

Note that the label limit (see Limitations, page 2-9)of 57,344 different labels includes labels from the bypassed VPNs.

Each bypassed VPN entry, both upstream and downstream, is removed from the database after a set period of time (10 minutes). If the entry is still used in the traffic, it will be re-learnt. This allows the database to remain clean, even if the labels are reused by the routers for different VPNs.

show bypassed VPNs In the **show bypassed VPNs** command, the age is indicated with each label - the length of time since it was learned.

Additional MPLS Pattern Support

The MPLS/VPN solution was designed to provide DPI services in MPLS/VPN network. These networks use BGP protocol as the control plane for the VPNs and LDP protocol for routing. There are complex networks where the MPLS infrastructure is used not only for VPN and routing, but also for other features such as traffic engineering (TE) and better fail-over. These features are usually enabled per VRF in the PE.

The Service Control MPLS/VPN solution does not support VPNs that use other MPLS-related features. Features such as MPLS-TE or MPLS-FRR (Fast Reroute) are not supported. VPNs for which these features are enabled can be automatically bypassed in the system, but are not allowed to be configured in the SM as serviced VPNs. Configuration of these VPNs in the SM might cause misclassification due to label aliasing.

The following list describes the labels combinations that are supported by the SCE platform and how each combination is interpreted by the platform:

- One or more labels upstream, no labels downstream:

  Assumed to be non-VPN (see Non-VPN-Based Subscribers, page 2-6).

  The SCE platform treats the following IP flows as non-VPN flows, and ignores their labels.

- One label upstream, one label downstream:

  Assumed to be VPN traffic, in which the P router happens to be the last hop in the upstream.

  The label in the downstream is treated as a BGP label, like the regular case. If the BGP label is known from the SM, then the flow is assigned to the correct subscriber, otherwise, it is treated as a bypassed VPN.

- Two labels upstream, one label downstream:

  This is the typical configuration of the system. Of the two upstream labels, one is for BGP and one for LDP. The downstream label is for BGP only

- More than two labels upstream, or more than one label downstream:

  These combinations occur when other MPLS-related features are enabled for the VPN. Such VPNs are not supported and should not be configured in the SM. However, they can be bypassed in the SCE platform without any service and without harming the service for other VPNs.

VPN Identifier (RD or RT)

Either the Route Distinguisher (RD) attribute or the Route Target (RT) attribute can be used to identify the VPN. It is required to decide which attribute best reflects the VPN partitioning, and configure the system accordingly. Note that the configuration is global for all the VPNs, that is, all VPNs must be identified by the same attribute.

The Route Distinguisher (RD) is generally used to distinguish the distinct VPN routes of separate customers who connect to the provider, so in most cases the RD is a good partition for the VPNs in the network. Since the RD is an identifier of the local VRF, and not the target VRF, it can be used to distinguish between VPNs that transfer information to a common central entity (for example a central bank, IRS, Port Authority, etc.).

The Route Target (RT) is used to define the destination VPN site. Though it is not intuitive to define the VPN based on its destination route, it might be easier in some cases. For example, if all the VPN sites that communicate to a central bank should be treated as a single subscriber, consider using the RT as the VPN identifier.

It is important to note that this configuration is global. Therefore, if at some point in time, any VPN would have to be defined by RD, then all the other VPNs must be defined by RD as well. This is a point to consider when designing the initial deployment.

Service Control MPLS/VPN Requirements

Topology

Following are the general topology requirements for MPLS/VPN support:

- The SCE platform is placed in the network between the P routers (Provider MPLS core) and the PE (Provider Edge) routers.
- The subscriber side of the SCE platform is connected toward the PE router.
- The network side of the SCE platform is connected toward the P router.
- The BGP LEG is installed on the SM, and is placed somewhere in the network.

  It speaks with the SCE platform through the management IP.

In a cascade installation:

- The two SCE platforms are connected to each other via the cascade interfaces.
- The data link between the P and the PE is connected via the other interfaces on each SCE platform, as described above:
  - Subscriber side of each SCE platform connected toward the PE router
  - Network side of each SCE platform connected toward the P router

The following drawing depicts a typical cascade installation.

Figure 2-1 Typical MPLs/VPN Installation

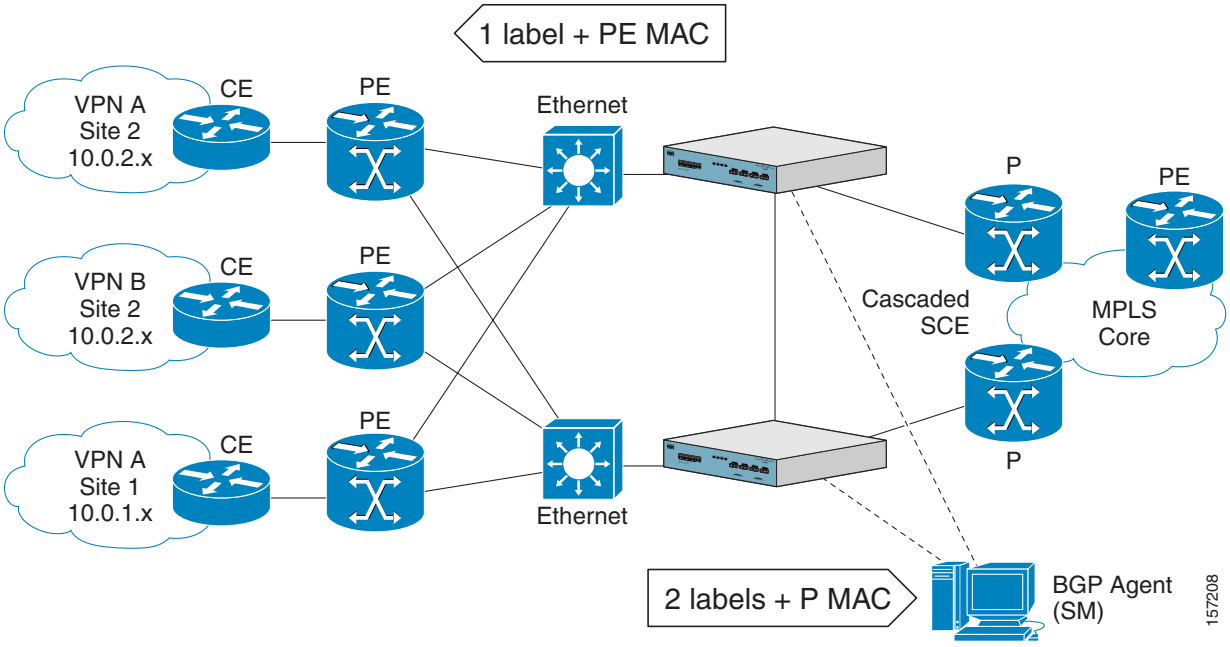

Capacity

The system supports:

- 2015 MPLS/VPNs
  - 80,000 IP mappings over VPNs
- 57,344 different labels (including upstream and downstream, and including the bypassed VPNs)
- 256 PEs per SCE platform
  - 4 interfaces per PE

Limitations

Note MPLS/VPN functionality is not supported on the Cisco SCE8000 platform.

Mutually exclusive system modes

When the system works in MPLS/VPN mode, the following modes are not supported:

- The following tunneling modes:
  - MPLS traffic engineering skip
  - MPLS VPN skip
  - L2TP skip
  - VLAN symmetric classify

- TCP Bypass-establishment
- DDoS
- Value Added Services (VAS) mode

Number of MPLS labels

- The choice of the unique VPN site must be based on the BGP label only. The BGP label must be the innermost label.
- The MPLS/VPN solution supports various combinations of labels. See Additional MPLS Pattern Support .
- The system does not support VPNs for which other MPLS-related features, such as MPLS-TE or MPLS-FRR, are enabled.

Subscriber-related limitations

The following subscriber-related limitations exist in the current solution:

- The SM must be configured to operate in Push mode.
- VLAN-based subscribers cannot be used.
- Introduced subscriber aging is not supported when using VPN-based subscribers.
- Maximum number of VPN-based mappings per single subscriber:
  - 200 (standalone)
  - 50 (cascade)

Topology-related limitations

- An asymmetrical routing topology in which the traffic may be unidirectional, is not supported, since the MPLS/VPN solution relies on the bidirectional nature of the traffic for various mechanisms.

TCP related requirements

- Number of Upstream TCP Flows – There must be enough TCP flows opening from the subscriber side on each PE-PE route in each period of time. The higher the rate of TCP flows from the subscriber side, the higher the accuracy of the mechanism can be.

VPN configuration requirements

- Two VPN sites must be aggregated into one VPN if the following conditions are both true:
  - They are both connected to the same SCE platform
  - They both communicate with a common remote site using the same upstream labels and P router.
- An MPLS/VPN-based subscriber MAY NOT have IP mappings over more than one VPN.

Backwards Compatibility

An SCE platform running SCOS V3.1.5 and up does not support MPLS/VPN subscribers of the type used in older versions. Instead of defining an MPLS/VPN subscriber, which reflects the whole VPN, the user must configure a VPN entity and a full range private IP subscriber within that VPN (0.0.0.0/0@VPN1)

When working with the combination of SM of a version before V3.1.5 and an SCE with V3.1.5 and up, only regular IP subscribers are supported. VPN-based subscribers are not supported at all in this combination.