

Selective Virtual Routing and Forwarding Table Download: A solution to increase Layer3 VPN scale

White Paper

What You Will Learn

This document considers the growing problem of the scale of Virtual Routing and Forwarding (VRF) tables and does the following:

- Provides an overview of Selective VRF download solution
- Introduces the filtering benefits of the solution
- Calculates savings and gains
- Presents some recommended deployment best practices

Background

With the advent of Layer 3 Virtual Private Networks (L3VPNs), whereby service providers can interconnect multiple sites of a customer network over a shared provider core or backbone network, the scale requirements for service provider routers are growing constantly. Addition of new customers for VPN services and increasing loads from existing subscribers are contributing to the growth of VRF tables. This has created challenges such as increased router loading times, longer maintenance windows, and frequent memory upgrades.

According to some projections, VPN route scale for some large service providers has reached 500,000 routes and is growing. With the global IPv4 Border Gateway Protocol (BGP) routing table occupying 300,000 routes, service providers and indirectly network equipment vendors are finding it increasingly difficult to scale the existing network infrastructure to support the increased demand.

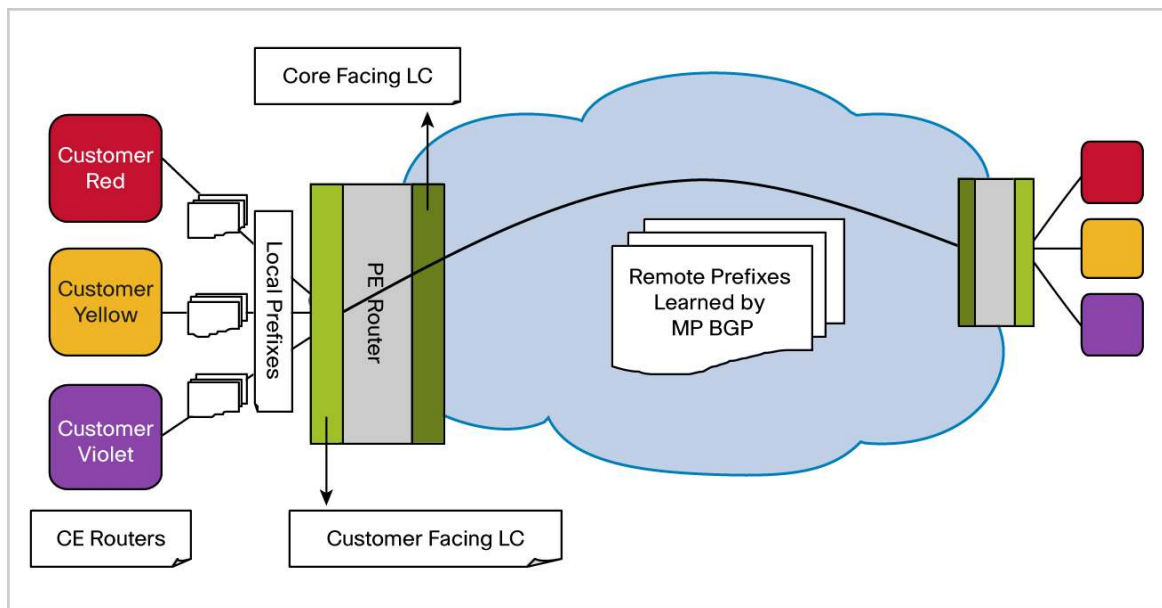
There has been much focused innovation on technologies to reduce VPN route scale based on positioning of the routers in the service provider network, including rt-constraint and ext-comm ORF. But these technologies approach optimization from the network design and control plane scale perspective. More can be done within the device: Selective VRF Download (SVD) provides a much-needed solution.

Solution Overview

Selective VRF Download is an architectural solution to alleviate scalability and convergence problems with L3VPN technologies. SVD is mainly applicable to service provider routers with multiple line cards (LCs) and forwarding engines. The basic idea is to download to a line card only those prefixes and labels that are actively required to forward traffic through that line card.

The SVD solution tries to optimize the table and route download to cards by classifying line cards by SVD card role and routes by SVD route type.

Figure 1. Selective VRF Download Basic Technology



The following are the SVD card roles:

- **Core facing role:** A card that has only core-facing interfaces (interfaces that connect to other provider core and provider edge devices). These are the interfaces which would be provisioned under default VRF.
- **Customer facing role:** A card that has only customer-facing interfaces (interfaces that connect to customer edge devices in named VRFs).

It is important to note that card roles are per address family.

The following are the SVD route types:

- **Local route:** A route that is received from a customer edge device connected to the router in a configured VRF context.
- **Remote route:** A route received from another provider edge device multiprotocol BGP (MP BGP) and is imported to a configured VRF.

Other important terms include:

- **Locally significant table:** When a VRF is provisioned to an IP numbered interface that has its data plane on the card, it will be considered as a locally significant table (LST) for the card. VRF tables that are LST for a card are downloaded to the card.
- **Inter-table dependency:** From the definition, VRF is an independent routing table instance. Under most instances a VRF will not be dependent on another instance, but there are two main exceptions where routes could have next-hops in another VRF:
 - Inter-VRF static routing
 - BGP extranets

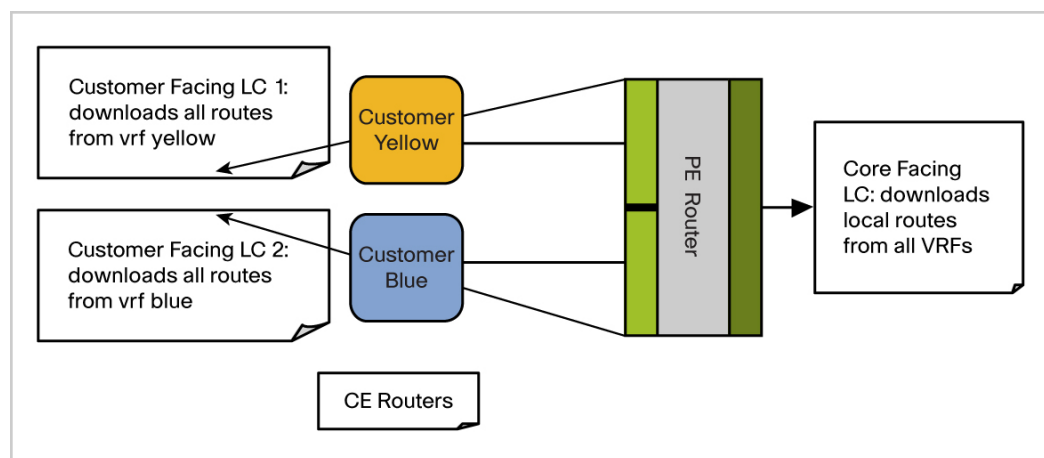
If a VRF table is dependent on another table, then that table also becomes LST (though indirect) and needs to be downloaded to the card. This process continues until all recursions are resolved, and it is performed on each card in a truly distributed manner.

VRF and Route Filtering

VRF and route filtering are performed based on the SVD card type:

- Core facing cards download routes for all VRFs, but only the local routes. Route filtering is done on the forwarding cards based on card role and route type.
- Customer facing cards download all routes only for VRFs that are significant to the line card. Table filtering is done on the route processor card based on the LST download requests from the forwarding cards.

Figure 2. Selective VRF Download Filtering



Savings Calculation

In Table 1:

- n is the total number of VRFs present
- o is the number of VRF directly provisioned/configured on the card, ($n \geq o$)
- R is routes per VRF
- x is the ratio of SVD local: total routes
- Y is the number of VRFs dependant on directly provisioned VRFs (o), ($Y \geq 0$)

Table 1. Total Tables and Routes Downloaded by Line Card Type

Card Type	Tables Downloaded	Routes Downloaded
Customer	$(o+Y)$	$(o+Y)R$
Core	n	nxR
Without SVD	n	nR

The preceding table summarizes the total routes and tables downloaded on the line cards of each SVD card type. Savings can be calculated by the difference between the numbers in the Without SVD row.

For example, a customer has 100 VRFs configured on the system, with five line cards. For the IPv4 address family, four line cards are working as customer facing with equal VRF distribution, while one is core facing. Inter-table dependencies do not exist. In this example, $n = 100$, $o = 25$, $x = 3/10$, $Y = 0$, and $R = 1000$.

Number of routes downloaded:

- Without SVD: $(nR) = 100,000$
- On customer-facing card: $(o+Y)R = 25,000$
- On core-facing card: $(nR) = 30,000$

In this example, the SVD feature brings almost 70 percent savings.

Recommended Deployment

- Separate customer cards from core. Planning a specific role for each card, rather than keeping both global and customer interfaces on the same card, brings the maximum benefits and savings. We recommend not mixing cards for different SVD roles.
- Try to minimize the dependency between VRFs. Interdependence among tables causes more and more tables to be downloaded to each card. This recursion could be difficult to realize and break. Hence we recommend avoiding multilevel recursive dependency among VRFs.
- Keep your interface configuration clean. Stray configurations can cause more VRFs to be considered locally significant and be downloaded. Those VRFs could further have dependencies and those dependencies would also be downloaded.

SVD on Cisco IOS XR Software

Selective VRF download can be deployed on any routing device with multiple forwarding cards. This section covers some of the extensions and specifics of the Cisco® IOS XR solution. Please go to [Cisco® IOS XR release documentation](#) or contact your account team for feature availability information.

Extensions

- **Per address family filtering:** The SVD solution is IP address family independent. This means that a card role can be calculated or assigned for each address family. VRF and route filtering also occurs per address family.
- **Not-interested role:** If a card is not participating in forwarding for a particular address family, the card is classified as having a not-interested role and no tables or routes for the address family are downloaded.
- **Pure L3VPN role:** For pure L3VPN customer cards, where global routing information is not needed, default table BGP routes are filtered and only IGP routes in the default table along with VRF routes are programmed. This helps in further scaling the forwarding cards.
- **Automatic card role calculation:** The SVD feature uses interface-VRF mappings and interface-IP configuration to automatically predict the user intent for deployment of a card and need for a particular VRF on a particular card. Only if a VRF is provisioned to an IP-assigned interface will it be considered an LST. In certain deployment scenarios such as 6VPE, where routes from one address family recurse over routes from another address family, the card roles are adjusted automatically to the optimal roles.
- **Automatic SVD route type promotion:** The SVD feature automatically promotes the SVD route type of remote routes as local if there exists a recursion of another local route over it. This ensures that due to filtering of remote routes on Core facing cards, local routes recursing over remote routes are not left unresolved.

Specifics

- The default VRF is present on all cards. Cards with Not-Interested role for an address family create it but do not contain any routes in it for that address family.
- For an IPv4 address family, the card role never automatically arrives as not interested. This is done primarily because of the lack of pure IPv6 deployments.

Savings Metrics

Table 2 and Figures 3 and 4 show results from some of the convergence and memory gains tests performed under common L3VPN deployment profiles.

Table 2. System Profile used for savings measurement

Profile Information	
System Scale	
Number of VRFs	800
VPNv4 routes	300K
VPNv6 routes	40K
Customer Facing Card Scale	
Number of VRFs	400
VPNv4 routes	150K
VPNv6 routes	20K
Core Facing Card Scale	
Number of VRFs	800
VPNv4 routes	60K
VPNv6 routes	8K

Figure 3. Ternary Content Addressable Memory (TCAM) Savings





Figure 4. Convergence Gains



Conclusion

- The SVD feature optimizes the VRF and route download to multcard routing systems, thereby saving precious hardware memory and network processing unit resources. Existing low-memory forwarding cards can be deployed as core facing, thereby protecting your investment.
- SVD improves scalability, reduces convergence time, and does not affect existing features such as BGP Prefix-Independent Convergence (PIC). If VRF is considered LST for a card, only then it will be downloaded to the card. This helps ensure that only what is actively required is downloaded. VRF and route filtering are based on SVD card role.

For More Information

For more information, see the following:

- [RFC 4364, BGP/MPLS IP VPNs](#), Rosen and Rekhter, February 2006
- [RFC 2858, Multiprotocol Extensions for BGP4](#), Bates, Chandra, Katz, and Rekhter, June 2000
- [RFC 5291, Outbound Route Filtering Capability for BGP-4](#), Chen, and Rekhter, August 2008
- [RFC 4684, Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching\(BGP/MPLS\) Internet Protocol \(IP\) Virtual Private Networks\(VPNs\)](#), Marques, P. et. al, November 2006



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Printed in USA

C11-681649-00 08/11