Layer 3 Virtual Switching

Integrated Virtual Routing with Multi-Layer Switching

Abstract

As today's enterprise networks continue to converge, there is an increasing need for logical separation for applications, and/or user groups within the network. Virtualization abstracts the physical infrastructure creating one or more virtual (logical) instances running on a single physical resource. Layer 3 Virtual Switching allows partitioning of a single switch into many virtual routers. A virtual router has the same capabilities and properties as a physical router. It inherits all the same routing mechanisms for configuration, operation and troubleshooting. As a result, each virtual switch domain can be separately managed and isolated for security safety measures

In this paper we'll provide an overview of the Layer 3 Virtual Switching technology, including the key benefits. In addition, this paper will provide an example of a security application that demonstrates Layer 3 Virtual Switching value.

Introduction

In the early 1990s Multi-Protocol Routers provided the engine to handle packet processing in enterprise and service provider backbones. The breadth of functionality integrated into Multi-Protocol Routers was critical for handling the diverse and complex requirements of Layer 3 networks such as IP, IPX, SNA and Decnet running on different Layer 2 networks such as Token-Ring, Ethernet and FDDI. That "breadth" came at a very high cost for the users and performance could not keep up with the growth in network traffic.

In 1997, Extreme Networks® released its Summit® product line, and changed the networking industry. Extreme Networks saw that the industry was converging around the Ethernet and IP protocols, integrating Ethernet and, IP processing into fast, compact, and cost-effective ASICs. The Summit line could perform at 10X the performance of a Multi-Protocol Router at 1/10th the price. The era of Layer 3 Switching had arrived.

Fast forward to 2004. Today, Extreme Networks is innovating by integrating "Virtual Routing" technology with multi-layer switching for enterprise networks the same way it innovated by integrating Layer 3 routing with Layer 2 switching in 1997. Extreme Network has "virtualized" Layer 3 routing technology into its ExtremeXOS™ network operating system and its 4th Generation Networking Silicon System (4GNSS) ASICs for the BlackDiamond® 10808 making the BlackDiamond the first product to incorporate Layer 3 Virtual Switching in the market.

Recently, enterprise and service provider networks have been demanding more flexibility, partitioning and security in the design and operations of their networks.

VLAN technology alone has limitations that fail to allow it to meet these emerging needs. Layer 3 Virtual Switching technology from Extreme Networks is designed to complement VLANs by providing additional layers of separation, control and security for network designs and operations. As a result, Layer 3 Virtual Switching delivers higher levels of availability and security than what is possible with VLAN technology alone.

Layer 3 Virtual Switching helps enable network managers to deal with rapid changes in their business requirements such as accommodating a new organizational structure and changes in their network requirements such as dealing with unpredictable network traffic patterns.

Virtual Switching Concepts

The concept of resource virtualization is not new to the IT industry. Virtualization of system software and hardware resources has been applied to various IT infrastructure and applications related to data center computing and storage.

By introducing Layer 3 Virtual Switching, Extreme Networks brings the concept of virtualization to Multi-Layer Switching. Layer 3 Virtual Switching implements a number of virtual routers (VR) on the same physical switch. A VR has the same capabilities and properties as a physical router does. It inherits all the same routing mechanisms for configuration, operation and troubleshooting. In a nutshell, there are many instances of a router and protocol code running on an ExtremeXOS switch, each running independently from one another and for a different purpose.

Layer 3 Virtual Switching is the perfect solution for partitioning a network and consolidating multiple Layer 3 routing domains into a single switch. Its benefits are significant since it reduces capital and operations costs for equipment, space, power and management.

Figure 1 shows the basic concept of Layer 3 Virtual Switching technology. The figure depicts three VRs implemented within a BlackDiamond 10808 switch where each VR is used to partition a large network.

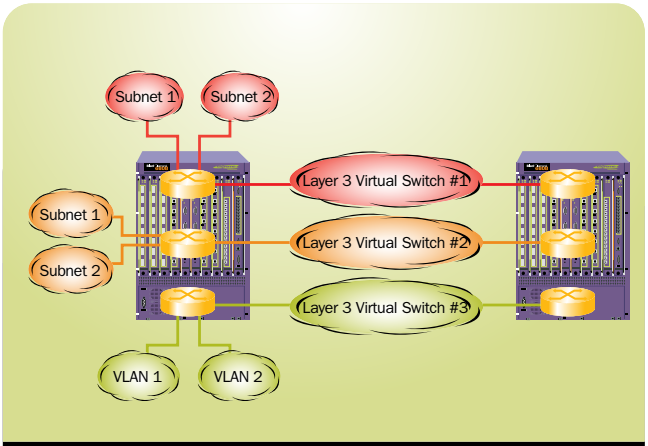

Figure 1: Extreme Networks Virtual Switching

Each VR in a Layer 3 switch that supports Layer 3 Virtual Switching technology maintains a separate logical forwarding table, that allows the VRs to have overlapping IP address spaces per VLAN. Since each VR maintains its own separate routing information, and switch ports can belong to only one VR, packets arriving at a port on one VR can never be switched to the ports on another. PIM, RIP, OSPF or BGP routing protocols can be used for each VR.

Figure 2 shows how Layer 3 Virtual Switching technology can be extended to multiple physical switches. The VRs with identical names are configured the same way in both the switches. Physically connecting the VRs with the same names together allows the Virtual Routing domains to be extended to both the switches.

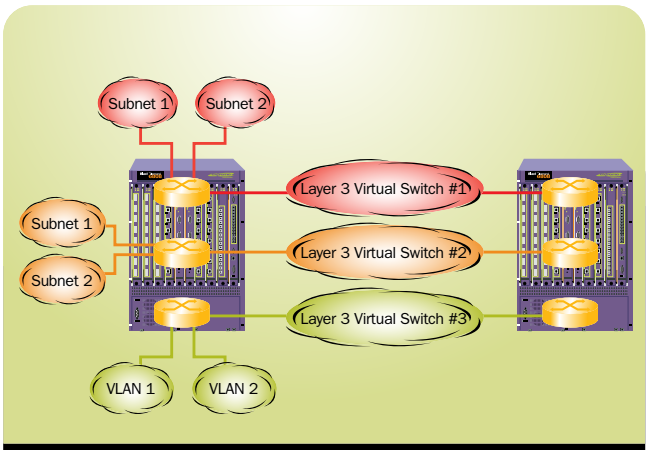

Figure 2: Extreme Networks Virtual Switching

Figure 3 gives an example of linking different administrative domains to different parts of the organization. In this scenario the organization is divided into three routing domains: one for the marketing and finance organization, one for the various research and development (R&D) labs and one for the company's extranet.

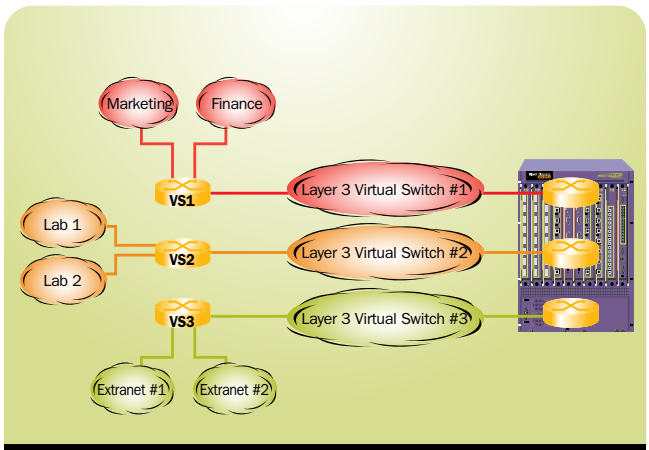

Figure 3: Partitioning a Corporate Network

Layer 3 Virtual Switching Applications

Network Segmentation: Different Administrative Domains

In a large company, different network managers have responsibility for different aspects of network operations. In a typical Layer 3 enterprise switch, access is "all-or-none." There is no way to administratively segregate certain ports or certain VLANs for one group and another set of ports and VLANs for another group.

Until now, one of three solutions have generally been adopted:

1. Have multiple organizations purchase and control their own switches.
2. Make one organization responsible for configuration, with "secondary" organizations putting in formal requests to have their changes made.
3. Allow both organizations to make changes, and trust them to be "cautious."

In practice, (1) results in higher costs, whereas (2) and (3) generally result in frustration or stability issues. A much better answer would be to separate Virtual Routing domains into separate administrative domains that could be independently configured by the separate IT organizations.

There are many other examples of an organization having the need to break a physical network into different administrative domains:

- In a converged enterprise, the VoIP network might be run by a different group than the data network
- In a large multi-site organization, the WAN network might be separate from the LAN network
- In a university, the computer science faculty's network might be a different administrative domain than the rest of the university

All of these are excellent applications for Layer 3 Virtual Switching, eliminating the need for each organization to purchase their own separate equipment.

Network Security: Trusted and Untrusted Domains

Best security practices have long called for IT operators to segregate their networks into "trusted" and "untrusted" domains. Figure 4 shows the concept of the trusted and untrusted domains. In the trusted area, users can get access to company servers, e-mail, and other resources. In the untrusted area, there are two options: purgatory and basic Internet access in "guest" mode.

Option 1: Purgatory

The first option for the untrusted network is to treat it as "purgatory." While in purgatory, users cannot access any internal or external resources whatsoever. The reason to have the untrusted domain is so that the user who has been denied access can (a) access a server that tells him why he has been denied access and (b) perform some action that will allow him to pass the screening process.

Example:
If a user does not have up-to-date virus definitions, then the user will be placed in the untrusted domain. In this state, the user may have access to a server from which the latest virus definitions and the latest version of the virus detection software can be downloaded. Additional steps like running a virus scan with the latest software and virus definition may also be initiated.

Option 2: Basic Internet Access in "Guest" Mode

Instead of purgatory, the organization might choose to admit unauthorized users to a "guest" mode, where they get basic Internet access service but are denied access to sensitive company resources.

This mode is gaining popularity in many corporations, particularly as they deploy WiFi services. Visitors to the company may wish to use WiFi to access the Internet, but these users must not have any access to company servers, intranet, etc. The obvious way to solve this problem is to use 802.1x to authenticate all users. If 802.1x authentication fails for any reason, the user gets dropped into the "guest" or "untrusted" domain, and can access the Internet only. If 802.1x authentication is successful, the user gets dropped into the trusted domain and can access all resources.

Whether Option 1 or Option 2 is implemented, using Layer 3 Virtual Switching technology is the ideal way to implement the network. Untrusted hosts are completely firewalled from the rest of the network.

Network Design: Managing Network Instability

Network instability can be caused by unknown device interactions that occur when changes are made in the network. Configuration changes might result in route flapping, Spanning Tree instability, or other issues that destabilize the entire network. The more the network changes, the more likely it is that some unexpected interaction will occur that leads to instability

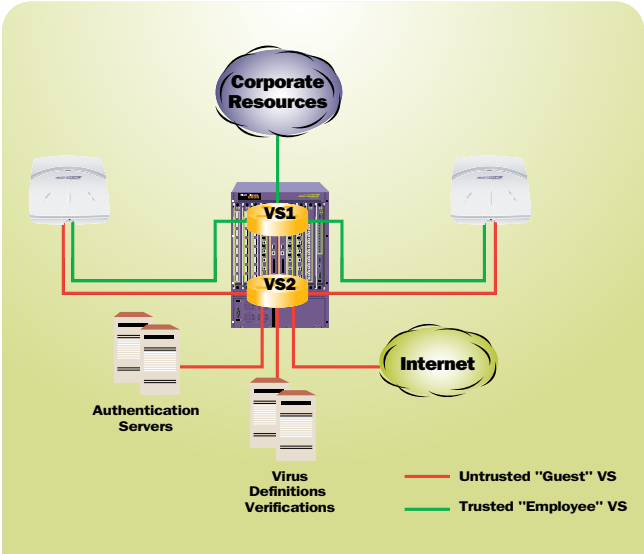

Figure 4: Trusted and Untrusted Domains

For example R&D labs have traditionally purchased and configured their own equipment—completely separate from the corporate network due to the potential for change related instabilities within the lab environment. For cost and efficiency, this is clearly not the best answer. A way of logically separating the "experimental" networks from the "mission-critical" networks so that they cannot impact one another is the right approach to solve this problem. Layer 3 Virtual Switching within a switch allows the network manager to do so. Figure 5 shows how an R&D Network and a Corporate Network can be separated using Layer 3 Virtual Switching.

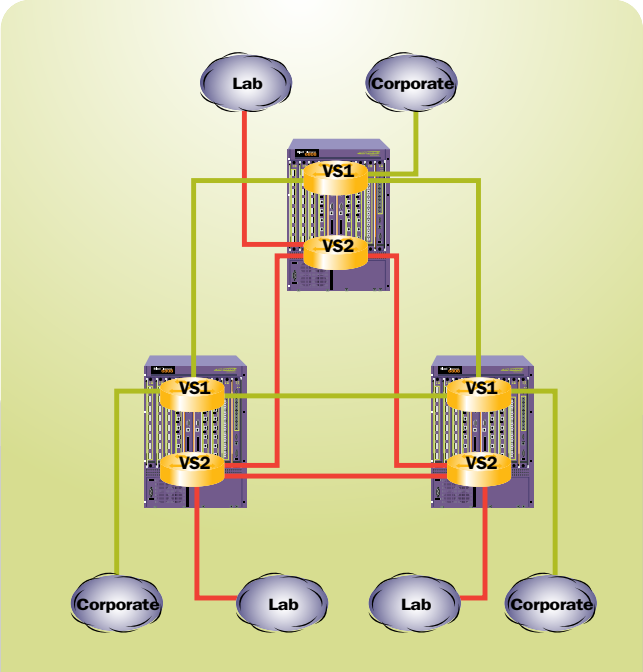

Figure 5: "Lab" and "Corporate" Administrative Domains

ExtremeXOS Virtual Routing Features

As outlined earlier, Layer 3 Virtual Switching is the integration of virtual routing with multi-layer switching. The following section describes the present implementation of virtual routing within Extreme Networks BlackDiamond 10808 and BlackDiamond 12804:

- Routed traffic is isolated in each VR within the switch. The user needs to externally implement cable connectivity between VRs to allow inter-VR traffic.
- Route table is separated per VR and overlapping routes are allowed between different VR.
- A VLAN can only be part of one VR. VLAN ID must be unique across the switch. Overlapping VLAN space and overlapping MAC addresses are not supported at this time. The system can support maximum 4,094 user-creatable VLANs. VLAN IDs are shared among all VRs. The user cannot pre-configure VLAN ID range to use for different VRs.
- Each VR will be assigned a set of ports. VLAN interfaces created in the VR can only contain ports assigned to the VR. Ports cannot be shared between VRs.
- The system resources of the switch are shared among all VRs. Resource over subscription is not supported. In other words, the resource consumed by all VRs added together cannot exceed system resource limitation. For example, if the system has a 20K ARP table entries limitation, then the ARP entries for all VRs added together cannot exceed 20K.

VR Types

There are two types of VRs in ExtremeXOS: System VRs and User VRs.

System VRs

System VRs are the special VRs used during system boot-up. These cannot be deleted or renamed. There are a total of three system VRs: they are VR-Management (VR-Mgmt), VR-Control, and VR-Default.

VR-Mgmt enables remote management stations to access the switch through Telnet, SSH, and SNMP sessions. VR-Mgmt owns the management port. No other ports can be added to the VR-Mgmt, and the management port cannot be removed from it. The management VLAN is created in the VR-Mgmt during boot-up. No other VLAN can be created in this VR, and the management VLAN cannot be deleted from it. No routing protocol will be running on this VR nor can any be added to this VR.

VR-Control is used for internal communications between all the modules and subsystems in the switch. It has no external visible ports, and no port can be assigned to it. VR-Control has no VLAN interface, and no VLAN can be created for it. No routing protocol can run or be added to this VR.

VR-Default is the default VR created by ExtremeXOS. All data ports in the switch are assigned to this VR by default. Any data port can be added to and deleted from this VR. Users can create and delete VLANs in this VR. The default VLAN is created in this VR during system boot-up. The default VLAN cannot be deleted from this VR. One instance of each routing protocol is started for this VR during boot-up, and these routing instances cannot be deleted.

User VRs

Users can create user VRs in addition to the system VRs. The total number of user VRs supported on the BlackDiamond 10808 and BlackDiamond 12804 with ExtremeXOS is 8. When a new user VR is created, by default, no ports are assigned and no VLAN interface is created.

Configuring Virtual Routers

Configuring VRs requires the following steps:

- Create the VR
- Add ports to the VR
- Add any required routing protocols to the VR
- Configure the routing protocols and VLANs

The following sections describe how to do these tasks.

Creating VRs

To create a user VR, issue the following command:

```
create virtual-router <vr-name>
```

A VR name cannot be the same as a VLAN name. User VRs cannot be named VR-Mgmt, VR-Control, or VR-Default because these names are reserved for the system VRs.

To delete a user VR, issue the following command:

```
delete virtual-router <vr-name>
```

Before you delete a VR, you must delete all VLANs created in that VR. All of the ports assigned to this VR will be deleted and made available for assignment to other VRs. Any routing protocol running on the VR will be shut down and deleted gracefully.

Adding Ports to a Virtual Router

By default, all the user data ports belong to the VR-Default, and belong to the default VLAN. A port cannot belong to more than one VR, so before a port can be added to a VR it may have to be deleted from another VR. As mentioned before a port must be deleted from any VLAN it belongs to before deleting it from a VR.

To add a port to a VR, use the following command:

```
configure vr <vr-name> add ports
<portlist>
```

To delete a port from a VR, issue the following command:

```
configure vr <vr-name> delete ports
<portlist>
```

The following is an example of removing all the ports on slot 3 from the default VLAN in the default VR and adding them to the VR named helix:

```
configure vlan default delete ports 3:*
configure vr vr-default delete ports 3:*
configure vr helix add ports 3:*
```

Adding Routing Protocols to a VR

Unlike the default system VR, VR-Default, no resources are allocated for routing protocols when a user VR is created. Routing protocols needed for a user VR must be added to the user VR before the protocols can be configured. When a protocol is added to a user VR, a process is started to support the protocol. Adding a protocol to a VR does not enable that protocol. You must then specifically enable and configure any protocol that you add.

To add a protocol to a VR, use the following command:

```
configure vr <vr-name> add protocol
<protocol-name>
```

To remove a protocol from a VR, use the following command:

```
configure vr <vr-name> delete protocol
<protocol-name>
```

Displaying Ports and Protocols

To display the ports, protocols, and the name of the protocol processes for a VR by using the following command:

```
show virtual-router {<vr-name>}
```

Configuring the Routing Protocols and VLANs

A user VR can be configured once the VR is created, the ports are added, and support for any needed routing protocols is added. To simplify configuring the user VRs, the concept of a VR configuration domain was added (instead of adding a VR keyword to every command in every routing protocol). VR commands are applied to the current configuration domain. The VR commands consist of all the BGP, OSPF, PIM and RIP commands, as well as the create VLAN and delete VLAN commands. Other commands apply to the switch as a whole.

To enter a VR configuration domain, use the following command:

```
virtual-router {<vr-name>}
```

For example, to enter the configuration domain for the VR helix, the Command Line Interface (CLI) session would look similar to this:

```
* BD10K.13 # virtual-router helix
* (vr helix) BD10K.14 #
```

The CLI prompt displays the VR configuration domain.

To return to the default configuration domain use the following command:

```
Use the VR command with no VR name, or
use the name VR-Default.
virtual-router
virtual-router VR-Default
```

To create VLANs in a VR, use the following command:

```
create vlan <vlan_name> {vr <vr-name>}
```

If you do not specify a VR in the create vlan command, the VLAN is created in the VR of the current configuration domain. The delete vlan command is also aware of the VR configuration domain.

All VLAN names and VLAN IDs on a switch must be unique, regardless of the VR in which they are created. You cannot have two VLANs with the same name, even if they are in different VRs. Routing protocols can be configured using the standard ExtremeXOS commands. The routing configurations of the different VRs are independent of each other.

Virtual Router Configuration Example

Following is a simple example of a VR configuration:

- The user VR helix is created
- Ports are removed from the VLAN Default and the VR VR-Default
- Ports are added to the VR helix
- OSPF is added to the VR helix
- The configuration domain is set to helix, so that subsequent VR commands affect the VR helix
- The VLAN helix-accounting is created
- Ports that belong to the VR helix are added to the VLAN helix-accounting

The CLI prompt is shown in this example to show how the VR configuration domain is displayed. At the end of the example, the VR is ready to be configured for OSPF:

```
* BD10K.1 # create virtual-router helix
* BD10K.2 # configure vlan default delete
  ports 3:*
* BD10K.3 # configure vr vr-default delete
  ports 3:*
* BD10K.4 # configure vr helix add
  ports 3:*
* BD10K.5 # configure vr helix add proto
  col ospf
* BD10K.6 # virtual-router helix
* (vr helix) BD10K.7 # create vlan helix-
  accounting
* (vr helix) BD10K.8 # configure helix-
  accounting add ports 3:1
* (vr helix) BD10K.9 #
```

Conclusion

IT managers today face tremendous challenges with network security and availability always rank at the top of the list. Layer 3 Virtual Switching integrates Virtual Routing technology with Multi-Layer Switching and gives IT staff a new tool for building a secure, highly available, and simpler network.

Implementing Layer 3 Virtual Switching is not an afterthought—hardware and software should be explicitly architected to support it. Designed for the secure and highly available enterprise, the 4GNSS chipset from Extreme Networks, along with ExtremeXOS, were designed from their inception with this capability in mind, making the BlackDiamond 10808 the first switch to support Layer 3 Virtual Switching in the industry.

www.extremenetworks.com email: info@extremenetworks.com

Corporate and North America
Extreme Networks, Inc.
3585 Monroe Street,
Santa Clara, CA 95051 USA
Phone +1 408 579 2800

Europe, Middle East, Africa and South America
Phone +31 30 800 5100

Asia Pacific
Phone +852 2517 1123

Japan
Phone +81 3 5842 4011