WHITE PAPER

Virtual Switches: Toward the Fully Secure Enterprise

extreme networks

Abstract - In the 1990's, VLAN's and Layer 3 switching technology transformed networking. These technologies are optimized to simplify and ease connectivity between all devices while allowing networks to scale gracefully.

More recently, emerging requirements for secure partitioning of the network. VLAN technology alone has limitations that fail to allow it to meet these emerging security needs.

Virtual Switching, a new technology proposed in this paper, is designed to augment VLANs by providing additional layers of security and control for network designs. Virtual Switching allow overlapping IP address space and VLAN space, independent administration and management capabilities, and resource separation. As a result, Virtual Switching delivers higher levels of security and availability than what is possible with VLAN technology alone.

Introduction

In the early 1990's, CPU-based multi-protocol routers provided the packet processing power in the network backbone. Multi-protocol routers' high degree of programmability was critical for handling the diverse and complex feature requirements of multi-protocol networks. But that programmability came at a high cost, and performance was limited.

In 1997, Extreme Networks released its Summit product line, and changed the networking industry. Extreme saw that the industry was converging around the Ethernet and IP protocols, integrating Ethernet and IP processing into fast, compact, and cost-effective ASIC's. The Summit line could perform at 10X the performance of a multi-protocol router at 1/10th the price. The era of the Layer 3 switch had arrived.

Fast forward to 2004. Today, Extreme Networks is doing for "virtual routing" what it did for routing and switching back in 1997 – moving capabilities into hardware. Extreme has integrated "Virtual Switch" (VS) technology the 4th Generation Networking Silicon System (4GNSS), with support from the new, modular ExtremeWare XOS operating system.

Virtual Switching enables a new level of security and availability – beyond that which could be achieved with the previous generation of Layer 3 switches. In today's Enterprise networks, when high levels of security are not a desire but a requirement, virtual switch capability is a must-have for any core switch deployment.

This paper will look in detail at how VS's can be used to build a secure, highly available Enterprise network.

Security in the Enterprise

Two 2003 surveys of IT managers[1] concur in that security is a top-ten issue for organizations.

There is good reason for security to be a top concern. Viruses, trojan horses, and worms are more and more frequent, and spread faster than ever before[2]. During the recent Slammer SQL worm outbreak, for example, the number of compromised hosts doubled every 8.5 seconds and was scanning 55 million IP addresses per second within minutes of its release.

The cost of these attacks is growing exponentially, and a single virus outbreak can now cost well in excess of $10 billion. What about total costs? In 2003, industry estimates pegged the total cost of virus attacks at around $55 billion, up from $13 billion in 2001 (cite source). (See Figure 1.) At the present rate of growth, virus damage will be more than $100 billion in 2004[3], and will rise to $1 trillion annually by 2008. IT managers feel the heat, and are trying to fight back.

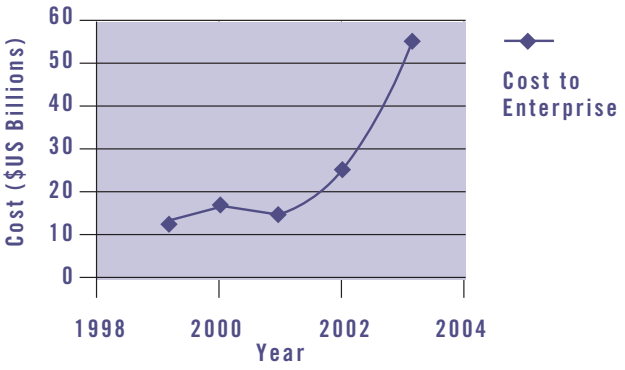

Figure 1. Cost of Virus Attacks

[1]Network World, "Survey: Security Cuts Both Ways on IP Plans," June 9, 2003
[2]????
[3]Computer World, Jan. 16, 2004

Virus attacks, while highly visible, are not the only security threats facing enterprises. Valuable intellectual property and financial data must be protected from unauthorized access and "hacker" attacks such as "man-in-the-middle" must be stopped.

How can the enterprise be safeguarded? Firewalls alone do not address the problem. While firewall technology is good at protecting the network perimeter, an estimated 80% of attacks originate inside the network (cite source). Installation of host-based virus protection software cannot address the issue either sinceeven the best virus protection does little good in the face of a new attack for which the software has no virus signature.

In fact, the entire network – from end to end – must be designed as a secure network. That means taking advantage of security features in enterprise switches. See Figure 2.

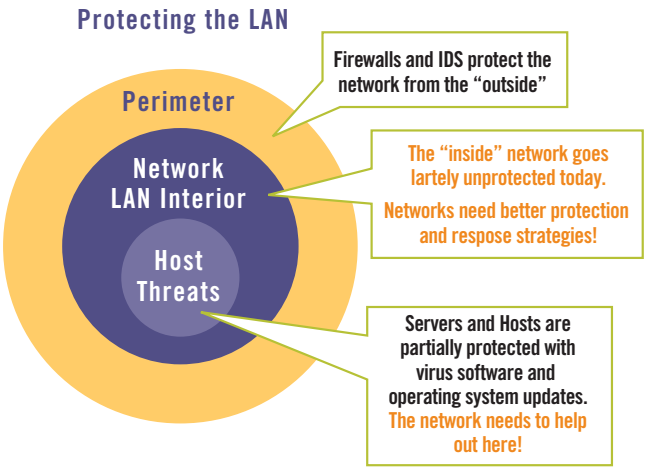

Figure 2. Protecting the LAN

Designing Secure Networks

VLAN's gained popularity in the 1990's as a means of grouping traffic for organizations. Particularly on wired networks, VLAN's could be associated with a particular port, protocol, or IP subnet, and security could be enforced; members of one VLAN could not talk to members of other VLAN's.

VLAN's, however, were only meant to optimize connectivity between disjoint groups; they were not meant to provide secure isolation between the groups. When used to implement security policy, a number of issues arise, including:

VLAN Leaking Issues

Some poorly designed switches allow, under certain circumstances, traffic to "leak" into other VLAN's, creating a major security hole. For this reason, many IT organizations refuse to use VLAN separate as the sole security barrier between untrusted domains.

Complex Policies in Routed Networks

Almost any VLAN enabled network will include routers to connect the different Layer 2 networks[4]. Complex ACL's or routing policies are required to stop traffic from being routed to or from "secure" VLAN's. Not only can ACL's and policies have a material performance impact on some networking equipment[5], but the additional complexity introduced with routing also increases the chances that human error will cause security to be compromised.

Limited Configuration Control

There is some danger that an operator will misconfigure the switch, and thereby allow traffic from different VLAN's to mix. Furthermore, there is no concept of having separate "VLAN domains" that can be controlled by different administrative entities, thereby limiting the ability to separate management domains to separate human administrators (further compromising security).

Lack of Resource Separation

VLAN's are ineffective in the face of DoS attacks because CPU and memory resources are all shared between VLAN's. An attack can quickly bring down the switch, thereby incapacitating all VLAN's running through the switch.

[4]Without routing between the Layer 2 networks, no connectivity would be available between the networks. Enterprise networks without routing enabled between the Layer 2 domains is almost unheard of.

[5]Cite Tolly report on Cisco 6500.

Availability in the Enterprise

In the network today, when more and more organizations are supporting mission-critical applications like VoIP on their networks, availability is more important than ever before. As switch capacity grows, so does the amount of data lost during the outage of any single networking component – and the cost associated with it. Figure 3 shows the cost per hour of a network outage across a number of different industries. In the financial industry, a network outage can cost up to $6.45 M per hour[6].

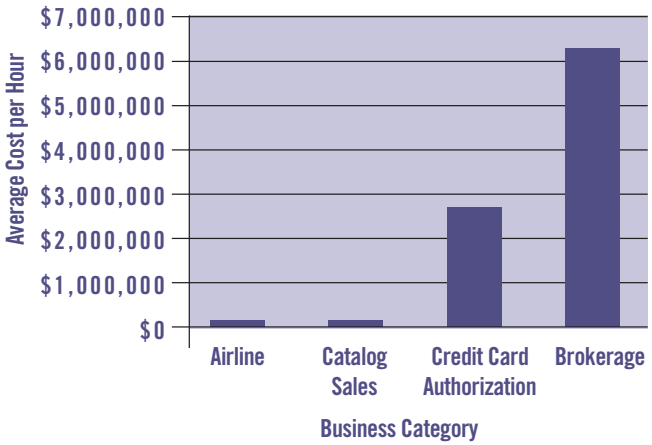

Figure 3. Average Cost of a Network Outage

Designing Available Networks

The obvious way for networking vendors to ensure maximum availability for their customers is building equipment with no single point of failure. On properly designed system, all components are redundant: control, switching, power supplies, fan trays, fan tray controllers, etc. If any one of these components fails, the system can stay up and running. A few core switches are starting to offer hitless failover and hitless availability features; if the control plane or switching fabric fails, the backup will take over without effecting any application traffic.

Unfortunately, eliminating single points of failure in the networking equipment only protects the network against hardware and software failures. In reality, there are other sources of potential downtime in current switch designs:

Operator Error: Complex Policy Configuration

One common source of operator error – tied in with the security concerns mentioned earlier – arises from the requirement to create complex policies for traffic control. Particularly in routed networks, complex ACLs or routing policies are required to stop traffic from being routed to or from "secure" VLAN's. These ACLs can effect performance, and the additional complexity increases the chances that human error will cause security to be compromised.

Operator Error: Lack of Administrative Domains

In a large company, different network managers will have responsibility for different aspects of network operations. For example, the IT department and its staff may have responsibility for the corporate network, while the Engineering group might have control over the R&D lab network. In a typical Layer 3 enterprise switch, access is "all-or-none." There is no way of administratively segregating certain ports or certain VLAN's for one group, other ports and VLAN's for another.

Until now, one of three solutions have generally been adopted: (a) have multiple organizations purchase and control their own switches, (b) make one organization becomes responsible for configuration, with "secondary" organizations putting in formal requests to have their changes made, or (c) allow both organizations to make changes, and trust them to be "cautious." In practice, (a) results in higher costs, whereas (b) and (c) generally results in frustration or stability issues. A much better answer would be the ability to separate ports and VLAN's into separate "administrative domains" that could be independently and securely configured by the separate IT organizations.

Network Instability: Frequent Network Changes

Instability can be caused by unknown device interactions that occur when changes are made in the network. For example, products from different companies may behave in unexpected ways when they are put in a network together. As well, software bugs in one piece of equipment can cause 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 and lead to instability.

It is because of this instability that R&D labs have traditionally purchased and configured their own equipment – completely separate from the corporate network. For cost and efficiency, this is clearly not the best answer. What is needed is a way of logically separating the "experimental" networks from the "mission-critical" networks so that they cannot impact one another.

[6]Dataquest report, "High-Availability Networking: Toward Zero Downtime," 2002

Virtual Switches

Virtual Switches (VS's) can complement traditional VLAN technology and provide additional security and availability options – addressing the security and availability limitations outlined in previous sections.

A single VS-capable switch can replace multiple physical switches; they do this by segregating internal physical resources (e.g., CPU, packet memory, and forwarding table space) and allocating them to the differentvirtual switches. These "virtual switches" can then be logically treated like separate switches, but each is housed in a single physical enclosure. See Figure 4.

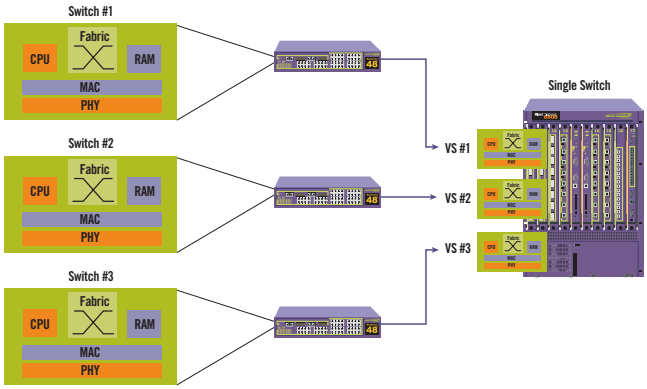

Figure 4. Virtual Switching

VS's differ from virtual routers (VR's) in being hardware-based implementations, with all ASIC's being VS aware.

Virtual Router (VR):
An emulation of a physical router at the software and hardware levels. Often characterized by the use of software-based routing technology, and typically used in the Service Provider space in conjunction with MPLS.

Virtual Switch (VS):
An emulation of a physical Layer 3 switch at the software and hardware levels. Characterized by ASIC-based forwarding that is "VS-aware," and capable of offering either Layer 2 bridged or Layer 3 routed services. Used in both Enterprise and Service Provider applications.

Extreme Networks has implemented a VS implementation, where all ASIC's are VS-aware. In Extreme's implementation:

- VS's have dedicated CPU cycles, forwarding table space, and packet memory on a per-instance basis, eliminating any possibility of leaking. This also serves to maximize availability, and prevent instabilities in one VS from affecting another.

- VS's dramatically simplify configuration by allowing the creation of separate administrative domains. Separate networks can have their own independent routing policies, eliminating the need to use ACL's and access policies. Different administrative entities can control their own VS, and access to other VS's will be blocked.

- VS's make the switch more resistant to attack, since resource separation means an attack on one VS will not bring down other VS's.

See Figure 5 for a logical picture of two switches, each with 2 VS's routing between multiple subnets, 1 VS running all Layer 2. Note that both L2 VS's and L3 VS's can be configured in the same physical switch. Here VS #1 might route traffic between Marketing and Finance; VS #2 might handle "Lab" traffic; and VS #3 might connect through a Firewall for an Extranet.

A sample code segment showing the configuration for VS1 is given below:

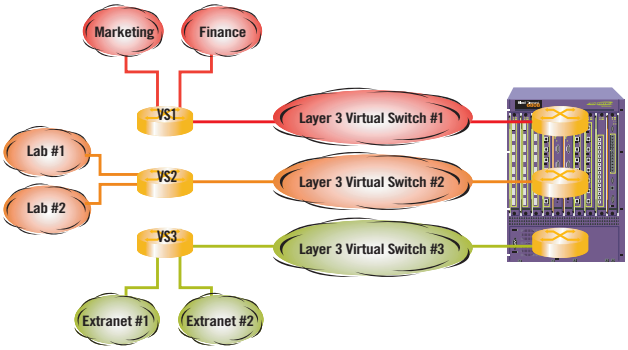

Figure 5. Layer 2 and Layer 3 Virtual Switches

```
#create virtual-switch vs1
#config vs1
#vs1> create vlan Marketing
#vs1> create vlan Finance
#vs1> create vlan InterSwitchVS1
#vs1> config Marketing ipaddress 192.168.0.1/24
#vs1> config Finance ipaddress 192.168.1.1/24
#vs1> config InterSwitchVS1 ipaddress 192.168.5.1/24
#vs1> enable ipforwarding
```

Enterprise VS Applications

This section will look at a number of common applications for virtual switches in today's Enterprise networks.

Application: Firewall

Many Enterprises would benefit from virtual switches in their firewall designs. Figure 6 shows a traditional firewall configuration. The Internet feeds into a "DMZ" or "De-Militarized Zone." The DMZ is outside the firewall and is therefore "exposed" to potential attack. Connected to the switches in the DMZ are all the servers that must be externally visible: WWW "Internet" servers, dial-up servers for the employees, VPN servers, etc. Redundant switches are often deployed, and servers are often dual-homed. In this way, failure of one switch will not result in loss of connectivity.

The switches in the DMZ generally do ACL filtering to protect the "exposed" DMZ servers from attack. e.g., WWW servers will only allow traffic bound for TCP port 80 (HTTP traffic).

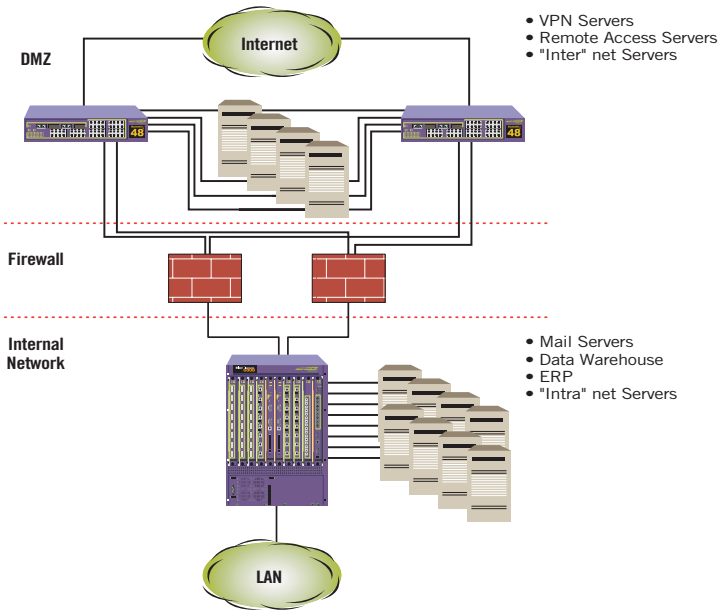

Figure 6.Traditional Firewall Architecture

The configuration shown in Figure 6 can be dramatically simplified with the help of virtual switches. See Figure 7. Rather than requiring three switches, with the use of virtual switches for the same results can be achieved with just one physical switch. Furthermore, there is no risk of a security breach or a Denial of Service (DoS) attack from the Internet affecting the VS to which internal LAN resources are connected. CPU, forwarding table, and packet memory is all separated on a per-VS basis. Even if a DoS attack were to bring down the VS connecting the DMZ, the LAN would be unaffected.

Depending on customer needs, a second switch may be used to provide additional fault tolerance. But if the core switch has high availability features such as those found in the BlackDiamond 10K, with redundant power, control, and switching, "Error Correction Code" (ECC) memory throughout the system, and a self-healing operating system with modularity, process restart, and memory protection – an additional switch is not necessary.

There are significant advantages to this architecture:

- Simple to manage, since there are fewer systems.
- Fully secure, with full segregation between VS's.

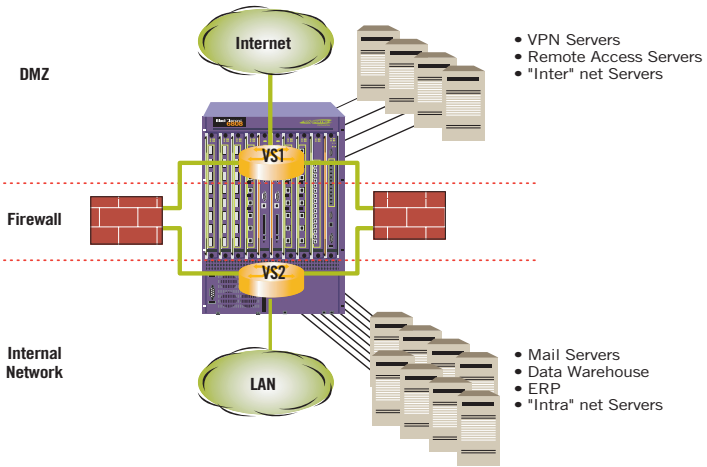

Figure 7. Firewall Architecture with Virtual Switches

- Can take advantage of more advanced services typically found on large core switch (e.g., ACL's; QoS capabilities and bandwidth guarantees; and CLEAR-Flow – which can be used to monitor traffic flows and dynamically shut down DoS attacks). (WEAK POINT)
- Easy to "extend" the DMZ into the LAN. This might be done if, for example, a lab server in a different building needed to be on the DMZ. In this case, the VS could be extended through the switches in the LAN – if they too had VS support – without risking LAN security.

Application: Extranets

In the firewall scenario, there were two virtual switches: one for the "trusted" internal network, and one for the "untrusted" Internet-accessible DMZ. Today, it is also very common for corporations to require a third VS for secure "partner" web sites, also known as "Extranets." In this scenario, a VPN connection is established between the

partner's autonomous system and the firewall. Traffic which has been authenticated is then permitted access to the "Extranet" web servers attached to the third virtual switch. See Figure 8.

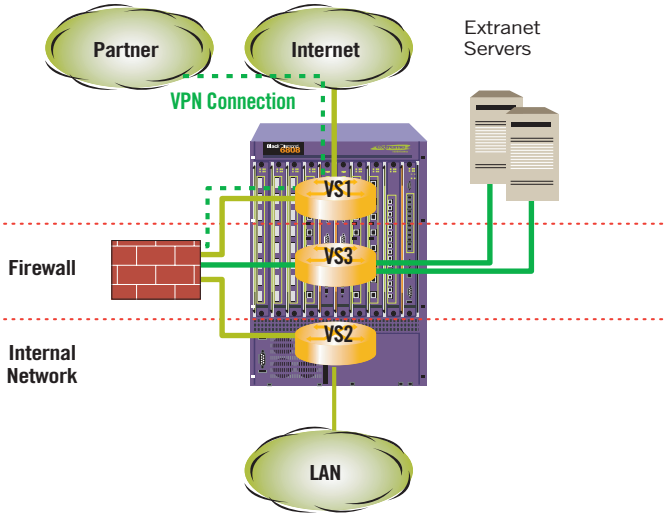

Figure 8. Extranets

Application: Administrative Domains

The firewall and extranet applications previously discussed illustrate the use of virtual switches for building up security and keeping a company's internal traffic separate from "external" traffic. But there are also good reasons to use virtual switches as a way to break up an enterprise's internal traffic into multiple "administrative domains."

Why would a company do this? The driving need is often organizational in nature. One typical example is a Research and Development (R&D) organization, where the IT department is responsible for running the corporate network, and the Engineering lab is run by the Engineering group. The IT group needs a rock-solid, stable network, whereas the Engineering group needs a flexible and adaptable network that can be changed at any time. In such a case, a good solution is to run the corporate network off one virtual switch, and the lab off another. The IT group and the Engineering group can each manage the resources under the control of their own virtual switch without any danger of disrupting the other group. See Figure 9.

There are many other examples of breaking the 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.

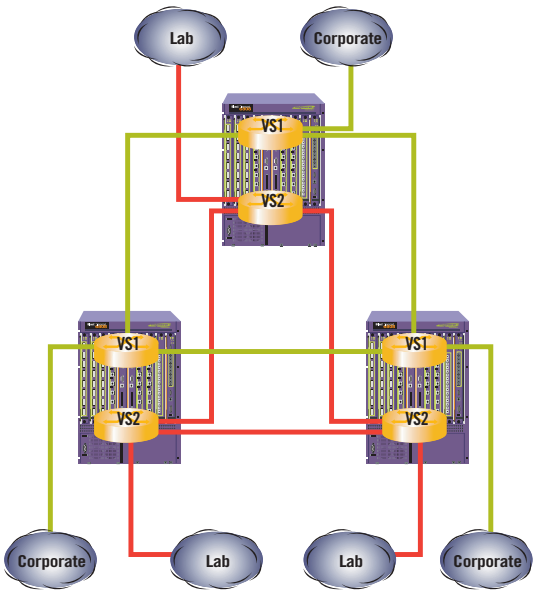

Figure 9. "Lab" & "Corporate" Administrative Domains

- 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 virtual switches, which eliminate the need for each organization to purchase their own separate equipment.

Application: Trusted and Untrusted Domains

Best security practices have long called for IT operators to segregate their networks into "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:

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, he will be placed in the untrusted domain. In this state, he may have access to a server which contains virus-detection software for installation.

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, Virtual Switch technology is the ideal way to implement the network. See Figure 10. Untrusted hosts are completely firewalled from the rest of the network. See Figure 9.

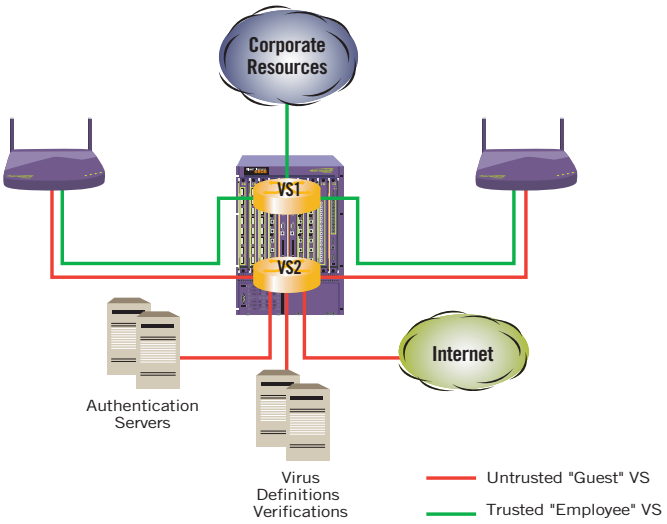

Figure 10. Trusted and Untrusted Domains

Requirements of Virtual Switches

Previous sections looked at why virtual switches are required, and how they are typically used. This section will look at what is needed to build virtual switch capabilities into a Layer 3 switch.

Separate Routing and VLAN Tables

At Layer 3, each VS instance must have its own independent routing table. Multiple organizations may use the same "private" IP addresses such as those assigned by RFC 1918 (10.x.y.z or 192.168.x.y). If this is the case, separate routing tables and virtual switches are necessary.

Beyond overlapping address spaces, another reason for having separate routing tables is so that routing instabilities or configuration changes in one "routing domain" cannot impact another. R&D organizations commonly fall into this category. In a company like Microsoft, for example, the "corporate" part of the network will generally be separated from the "lab" network. The lab is managed as a separate entity so that the potentially unstable lab network cannot bring down the entire corporate network. See Figure 11.

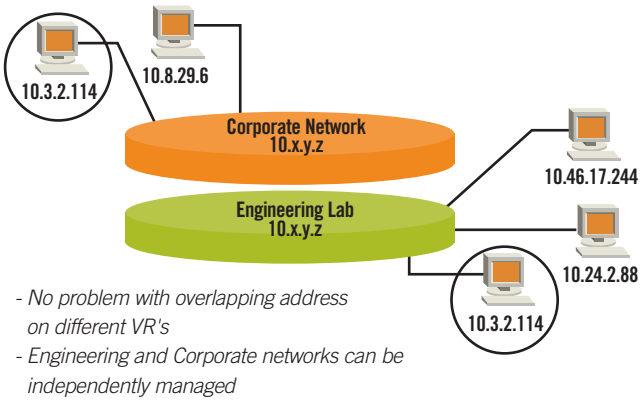

Figure 11. Separate Routing Tables in VS's

In a L3 VS, there must also be an ability to limit the forwarding table space and other resources that can be used by each VS. Without such safeguards, a single virtual switch could "hog" all the forwarding table memory, and leave insufficient CPU resources for other VS's.

Separate Management Instances

Standard MIB's fail to accommodate for the notion of Virtual Switching. For a Virtual Switch to truly look like separate physical switches, the Management Information Base (MIB) on each switch/router must be partitioned by VS.

Even on a single port, there may be multiple VS's connected. (Different VLAN's on a single port could be part of different virtual switches. For example, if there is a single fiber connecting two buildings, with R&D labs and corporate resources in both buildings, traffic from the R&D VS and the Corporate VS must use the same physical link.) In such cases, MIB II statistics such as bytes sent/received, packets sent/received, etc., must be collected for each VS independently as if they are on different ports, and then stored in different MIB objects.

In addition, there must be multiple "MIB views." The "owner" of a single virtual switch instance must be denied access to the MIB's of other virtual switches. In fact, the "owner" should not even know of the existence of other VS's or their MIB objects, since the whole objective is to give the appearance that a VS is a standalone switch.

Partitioned CPU and Memory Resources

Each virtual switch must be assigned a limited share of resources on the physical switch. Without limitations on CPU resources, a Denial of Service (DoS) attack on one VS has the potential to bring down every single VS on the switch. Similarly, a mis-configuration causing a "loop" on one VS must not affect others. Configuring and enforcing limits on the routing tables, VLAN tables and packet memory usage of the various VS's prevents them from starving each other.

Figure 12 shows the VLAN scenario, where a successful DoS attack on one VLAN brings down every VLAN on the switch.

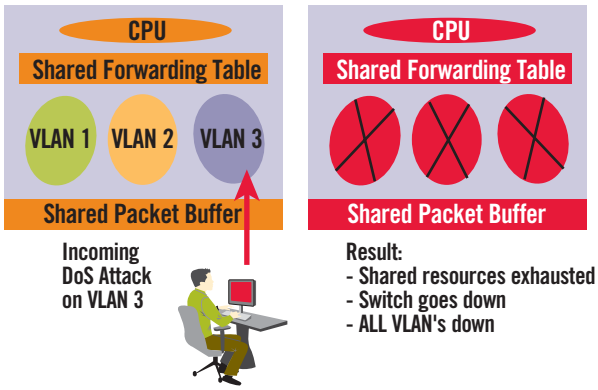

Figure 12. Lack of VLAN Resource Separation

Figure 13 shows what happens on a Layer 3 switch that has a VS implementation. Now each VS has dedicated CPU, buffer, and forwarding table memory. A successful DoS attack on one VS will only bring down that VS – and all other VS's will remain unaffected.

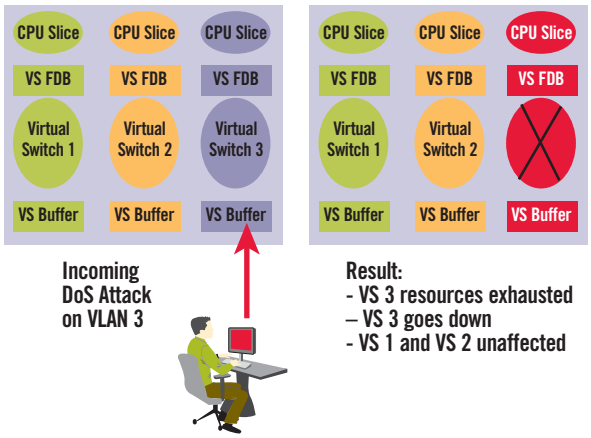

Figure 13. Resource Separation in VS Architecture

QoS per Virtual Switch

A requirement of virtual switches is that Quality of Service (QoS) guarantees be offered across shared links, and each VS be configured on independent queues – each with their own bandwidth guarantees. Otherwise, it would be possible for heavy traffic from one VS to "starve" another VS. Recall that the objective is to give the illusion that one VS is a physically separate switch; this breaks down if traffic from other VS's affect performance.

To support independent QoS across the VS, BlackDiamond 10K was built to support up to 128 queues per port – each of which can be independently assigned a minimum and maximum bandwidth. These can each be assigned to a VS.

Multiple Services per Virtual Switch

In order to give the full appearance of a physically separate router, each VS must offer the standard set of features available in a router. These features should include:

- BGP-4 Routing Protocol
- OSPF Routing Protocol
- Static Route Configuration
- PIM-SM / PIM-DM
- Layer 2 Resiliency Protocols
- Access Control List configuration
- QoS control, with ingress and egress rate shaping

Implementation Challenges

There are a number of challenges to implementing VS's in today's switches. On the hardware side, ASIC's must be designed with architectural support for virtual switching, providing segregation between different forwarding tables to ensure that overlapping address spaces can be supported. Ideally, this segregation is not solely controlled by indexing the tables by ingress port or VLAN. Truly flexible designs will use a separate VS-ID to index the tables, and will provide flexible mappings of ingress port, VLAN and ACL matching to determine the VS-ID to be used when performing looking for each individual packet. Virtually no chipsets support the capability today.

In addition to supporting VS's in the ASIC forwarding tables, virtual switch capable switches also require extensive CPU horsepower for managing the control plane. Each VS requires dedicated and distinct copies of control processes to manage tasks such as OSPF and Spanning Tree. Extreme's BlackDiamond 10K, designed

to support VS throughout the archtiecture, has extensible 4GNSS ASIC technology with dual 750MHz CPU's on each MSM, and 333MHz CPU's on each I/O module.

Software support for VS's is also non-trivial, and may be difficult or impossible to crowbar into an existing, monolithic operating system. Adding VS capability into an older O/S means fundamentally changing the routing stack at the heart of the operating system – a very difficult task.

ExtremeWare XOS was designed from the start with VS's in all of the data structures and VS support in all software modules. VS's are a fundamental part of the operating system, and represent just one of the advantages of ExtremeWare XOS.

Extensible Processing Power

As the number of virtual switches in a single switch increases, so does the amount of processing power required – particularly when each VS is running CPU-intensive dynamic routing protocols such as BGP-4. In ExtremeWare XOS, 10 Layer 3 VS's can run dynamic routing protocols. Up to 128 VS's can be supported if L3 VS's use static routes only.

ExtremeWare XOS is designed with a fully extensible control plane, allowing future expansion well beyond 10 VS's with dynamic routing. A layer of software called the "Inter Process Message Layer" enables additional CPU processing power to be added and leveraged from anywhere in the system. As a result, external CPU power (e.g., on a PC based platform) can be connected to a switch to augment VS processing. This capability provides an impressive level of flexibility that allows a single switch running ExtremeWare XOS to support hundreds of Layer 3 virtual switches – all running heavyweight dynamic protocols such as BGP-4.

Conclusion

IT managers today face tremendous challenges, and security and availability concerns always rank at the top of the list. Virtual Switch (VS) technology gives IT staff a new tool for building a secure, highly available network, a technology with many advantages over existing VLAN switching:

- Resource separation: An attack or network instability on one VS will not impact others.
- Administrative domains: Multiple groups can configure the same switch, and only access the VLAN's and ports in their own VS.
- VS-aware MIB's: Management and stats information only available to authorized personnel for each individual VS.
- Management simplicity: No need for complex policies to prevent L3 communication to or from secure VLAN's.

To support virtual switches, hardware and software should be explicitly architected to support it. Designed for the secure and highly available enterprise, Extreme's 4GNSS chipset, along with ExtremeWare XOS, were designed from their inception with VS capabilities in mind.

Appendix: Virtual Routers in the Service Provider Space

In the 1970's, when networking was in its infancy, corporations needed a way to connect up remote sites across different geographies. At the time the only option was to lease a private, dedicated line from their RBOC / PTT – a very expensive option, and one which did not make efficient use of the bandwidth capacity between locations. The charging model was not tied to the amount of traffic traversing the link.

As an alternative, Service Providers sought ways to use a common physical infrastructure, but to logically segregate individual customers on the same physical channel. By doing so, the cost of the physical infrastructure could be shared amongst multiple customers, thereby driving down cost per customer, and utilizing bandwidth more effectively. This approach was called a "Virtual Private Network" (VPN) because customers would enjoy a "leased line" experience, but would in fact be connected to a shared network.

The first effort to build VPN technology came with the X.25 protocol. X.25 had the concept of "Closed User Groups" (CUG's), which restricted data inter-connectivity to members of the CUG.

In the 1980's, first Frame Relay and then Asynchronous Transfer Mode (ATM) arose from the ashes of X.25, offering significantly more bandwidth and new capabilities. Each of these protocols could emulate a leased line by providing a "virtual circuit" between locations. A unique identifier inside the Frame Relay frame (the Data Link Connection Identifier, or DLCI) was used to switch traffic to the appropriate network destination, and only that destination. Frame Relay also

originated the concept of a bandwidth guarantee; Frame Relay users could be guaranteed a minimum level of bandwidth of "Committed Information Rate" (CIR) across the connection. ATM, which followed Frame Relay, used the VPI/VCI (Virtual Path Identifier / Virtuail Circuit Identifier) in place of a DLCI to deliver the same basic function. ATM also provided different levels of Quality of Service (QoS) that could be applied based on the QoS requirements of customer applications.

Changes in the LAN

As Frame Relay and ATM protocols were evolving in the Wide Area Network (WAN) during the 1990's, the Local Area Network (LAN) was taking a completely different course. In the early 1990's, a large number of different Layer 2 protocols (Token Ring, Ethernet, and then FDDI) and Layer 3 protocols (DecNet, IPX, Appletalk, IP) were vying for dominance. Through the end of the 1990's and into 2000, the simplicity and power of Ethernet and IP were rapidly making those two protocols dominant in the LAN.

A major part of this trend was the invention of the "ARPANET", a research network based on IP that later evolved into the Internet and "World Wide Web" that is ubiquitous today. LAN's began to use Ethernet and IP for their ease of connecting to the WWW and for their scalability as evidenced by the rapidly growing Internet.

Naturally, the increasing use of Ethernet, IP, and IP-based applications in the LAN led to increased bandwidth demands, and there was a need to connect multiple sites using VPN technologies. Frame Relay and ATM, both Layer 2 approaches, were used – but had a number of disadvantages:

- Inters-site connectivity required restrictive "Full Mesh" or "Hub-and-Spoke" architectures.
- Need for customers to understand and manage a complex IP network across multiple sites.
- Explosion in number of VC's prevented approach from scaling to hundreds of sites.

Virtual Routers

The "Layer 3 VPN" concept was developed as an answer to the disadvantages of the Layer 2 VPN model. The basic idea was to participate in the customer's Layer 3 routing topology, and use IP routing to direct traffic between sites. If all the Service Provider's boxes were IP-aware, it was no longer a requirement to use full mesh or hub-and-spoke topologies. The number of connections between sites was dramatically reduced.

Layer 3 VPN's use "virtual routers" to maintain separation between customers' VPN's; otherwise, there would be no separation between customer traffic, a key requirement of a VPN. A Virtual Router (VR) implementation in a switch / router emulates multiple physical routers in a single box. From the customer perspective, it appears that they have their own dedicated router network. Initially, VR implementations were mainly software-based, whereas an Extreme's ASIC's include VR support. The resulting hardware based virtual router is a subset of what Extreme Networks calls a "Layer 3 Virtual Switch." "Layer 2 Virtual Switches" are a parallel construct that provides separate, independent VLAN spaces. In fact, this is a Layer 3 Virtual Switch.

Service Provider Applications Today

Service Providers today make use of two types of VPN services: Layer 2 VPN's and Layer 3 VPN's. With a Layer 2 VPN, the service provider simply allows the enterprise to bridge information over a WAN or MAN transparently. It is as if the service provider is offering a large distributed bridge / Layer 2 switch. Common methods for offering Layer 2 VPN's include the "VMAN" (802.1Q-in-802.1Q) encapsulation approach or the Multi Protocol Label Switching (MPLS) technique with Hierarchicacal Virtual Private LAN Services (HVPLS).

The second type of VPN, the Layer 3 VPN, is more complicated. Here, rather than just transparently passing traffic between a customer's different sites, the service provider actually participates in routing traffic from site to site. Layer 3 VPN's, most commonly based on RFC 2547bis, are more complex and difficult to manage, but may scale better for very large VPN offerings because they eliminate the need for Layer 2 learning across a MAN / WAN and reduce the level of broadcast traffic.

Virtual routers are a requirement for offering Layer 3 VPN services. The whole idea of Layer 3 VPN services is to look (to the end customer) as if they have their own, private router. As such, the resources within the switch must be partitioned between the different virtual routers. If multiple virtual routers are supported, CPU time must be split between VR's. Each must have separate routing tables, separate instances of routing protocols, separate queues for SLA guarantees, etc.

The Layer 2 and Layer 3 VPN's used by service providers are complex to manage, and designed to support a very large number of VPN's that would be overkill in the Enterprise. It is for this reason that most Enterprises implement VS technology alone and prefer to avoid the complexities of protocols such as MPLS.

3585 Monroe Street Santa Clara, CA 95051-1450 Phone 408.579.2800 Fax 408.579.3000
Email info@extremenetworks.com Web www.extremenetworks.com