



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Introduction

One of the major shortfalls of a VPN (Virtual Private Network) solution in an enterprise network is the maximum bandwidth they can utilize. This is especially true when a VPN is used to connect remote sites to establish security. The VPN's performance is degraded when other security protocols are chosen. For example, a VPN that supports 100MB throughput, adding 3DES and SHA-1 could decrease its performance almost 50%. The VPN is not only the network's bottleneck, but also provides a single point of failure. For these reasons, many users choose not to implement a security solution such as VPNs. With the use of a loadbalancer, grouping multiple VPNs can achieve the type of bandwidth required to support high-speed connection lines and provide redundancy.

Purpose

This paper identifies the required sections needed for the configuration of the Foundry ServerIron XL to load balance the Alcatel VPN PERMIT/Gate. The VPN section only covers the information pertinent to the ServerIron XL. In the ServerIron XL section, it gives users an understanding of the required tasks needed to load balance traffic, provide fail-over, and take advantage of the maximum bandwidth available. In this example, the network will consist of three sites (Site A, Site B, and Site C) with two VPNs on each one. This solution offers the ability to scale. If another site is desired, a minimum of one ServerIron XL and two VPNs will be needed.

Product Notes

Although Alcatel has a complete family of VPNs, this paper focuses on the 7137 Secure VPN Gateway. The product is a tamper-resistant gateway that secures data communications for Intranets, Extranets, and Internet remote access. It supports various encryption protocols, authentication algorithms, and Public Key Infrastructure (PKI). Alcatel states that, "customers can seamlessly integrate the dedicated VPN network equipment devices with any existing PKI to conduct highly secure communications over the Internet between multiple sites."¹ Therefore the VPNs can be used in a network regardless of the PKI solution deployed.

The VPNs will be using IPsec (IP Security). According to an article published in the NetBSD web site, "IPsec provides per-packet authenticity/confidentiality guarantees between peers communicate using IPsec."⁴ IPsec is compliant in IPV4 and is a standard in IPV6. Communication can be accomplished using either transport mode or tunnel mode. In this configuration, IPsec will be using tunnel mode. For a complete tutorial on Ipsec, please refer to the Ipsec web site provided in the bibliography.

Some of the encryption protocols supported are DES, 3DES, and Blowfish. The algorithms are MD5 and SHA-1 group 1, 2, and 5. A brief explanation of these protocols will be discussed in this paper. When SHA-1 is selected, group 1 is the default. The user can select the encryption protocols and authentication algorithms through its security descriptor file.

3DES is a standard derived by DES. The National Institute of Standards and Technology (NIST) states that, "A DES key consists of 64 binary digits ("0"s or "1"s) of which 56 bits are randomly generated and used directly by the algorithm. The other 8 bits, which are not used by the algorithm, may be used for error detection."³ 3DES does exactly the same thing, only it runs the message through this algorithm three times.

W3C states that, "The Secure Hash Algorithm takes a message of less than 2^{64} bits in length and produces a 160-bit message digest which is designed so that it should be computationally expensive to find a text which matches a given hash."² For example, if userA sends a message to userB, the whole message is processed through the hash and the output produced is the same regardless of the size of the message. The hash, along with the unencrypted message is sent to userB. When the message is received by userB, the unencrypted message is processed through the same hash and the output is compared with the hash sent by userA. If the hash is not the same, authentication will fail and the packets are dropped. Group 1 refers to Diffie-Helman using a 768 key for the hash. Group 2 increases the key to 1024 and group 5 to 2048.

PERMIT/Gates can be configured in a cluster to restore a secure connection in the event of a Gate failure without data loss or service interruption. The VPNs in a cluster can be configured to authenticate to each other either through PKI certificates, or a shared password. A heartbeat is configured that specifies how often the secondary VPN will send a health check to the primary. Clustering has the following features:

- The PERMIT/Gate cluster appears to the outside world as a single gateway.
- The PERMIT/Gate cluster continues to function as long as at least one member gateway is operational.
- Recovery from failure is automatic.
- Individual gateways can join and leave the cluster without adversely affecting its operation.
- Inter-cluster communications are secure and tamper-proof.
- The PERMIT/Gate cluster can be reconfigured without shutting it down.

However, clustering uses only one VPN at any given time, therefore the bandwidth issue is not addressed. Although it provides great fail-over solution, the original problem exists. By placing a ServerIron XL in front of the VPNs, traffic is distributed evenly. The ServerIron XL can determine when a path is down and provide redundancy.

In this example, by having three sites, the VPNs and ServerIron XLs will be configured in a full mesh environment. If VPN1 on Site A fails, VPN1 on Site B and C will continue to operate. However, if a ServerIron XL fails on Site A, the whole site will be down. In this case, the ServerIron XL is the single point of failure.

The Foundry ServerIron XL is ideally a layer 2 switch used to load balance firewalls and other servers. When not configured, it acts like a hub. In this paper, the ServerIron XL will be configured for firewall load balancing, using VPNs. The ServerIron XLs is located on the red (unencrypted) side of the VPN.

VPN Setup

Although many configuration settings for the VPN will not be addressed, they are still required. In order for this setup to work, the VPN must be completely configured. The PERMIT/Gate needs two Internet Protocol (IP) addresses: one in the immediate local subnet on

the private, unencrypted (red) side of the PERMIT/Gate, and one in the immediate local subnet on the public, encrypted (black) side of the PERMIT/Gate. The ServerIron XL needs only the red side IP address.

Connect the cat 5 cable from the red port of VPN1 to the first port on the ServerIron XL. Connect the cat 5 cable from the red port of VPN2 to the second port on the ServerIron XL. A connection is required between the ServerIron XL and the internal router pointing to the local area network. Connect the black ports on both VPNs to the router pointing to the wide area network. The interface connected to the ServerIron XL, the ServerIron XL itself, and the red ports on VPN1 and VPN2 must be on the same subnet.

Secure Maps

The secure maps represent a static route that the VPNs must take to reach a specific network. The entries in the map are in essence, black port IP addresses of the distant VPNs. These maps are crucial in the successful operation of this setup. When creating the secure map, ensure the proper format is followed. Any changes to the format will prevent the VPN from establishing a tunnel.

The secure map requires a one-to-one map (e.g. VPN1 on site A must be mapped to VPN1 on site B; VPN2 on site A must be mapped to VPN2 on site B; and so on). IF VPN1 is misconfigured and the map is created with VPN1 and VPN2 on Site B, the tunnel will not be established with either one. When creating a secure map, the red network of the local VPN with its black interface is added by default; this is the first entry. Do not make any modifications to this entry. Enter an entry for the distant VPN.

The # sign comments the line (The VPN won't read anything that follows a # sign). The first line read is the version number. The begin static-map entry informs the VPN that a static entry follows. The target entry is the destination's red network, or the protected network. The mode statement identifies what the distant VPN will use for authentication. The tunnel entry identifies the black IP address of the VPN that protects the destination network. This is a sample secure map.

```
# PERMIT/Gate Secure Map
# Generated at: 5/May/2001 04:01:03PM
# Generated by: PERMIT/Config version 3.00.012
version 1
begin static-map
target "1.2.3.*"
mode "ISAKMP-Shared"
tunnel "5.6.7.8"
end

begin static-map
target "9.10.11.*"
mode "ISAKMP-Shared"
tunnel "13.14.15.16"
end
```

```
begin static-map
target "17.18.19.*"
mode "ISAKMP-Shared"
tunnel "21.22.23.24"
end
```

The configuration means that in order to get to the 1.2.3.* network on Site A through its VPN1, a tunnel will have to be established through its 5.6.7.8 Black Interface using a shared secret password. This entry will be automatically entered when the secure map is first created. Even though this entry is for your own network, do not delete it. In order to get to the 9.10.11.* network on Site B through its VPN1, a tunnel will have to be established through the 13.14.15.16 Black Interface using a shared secret password. In order to get to the 17.18.19.* network on Site C through its VPN1, a tunnel will have to be established through the 21.22.23.24 Black Interface using a shared secret password. Notice that there are no entries for VPNs 2.

Configuring the ServerIron XL

The ServerIron XL is configured using a command line interface similar to Cisco's IOS. Configuration is done through its serial port. The ServerIron needs an Internet Protocol (IP) address for management purposes and so that it can be pinged from its peer ServerIron XL. The IP address must be in the same subnet as the red ports of the VPNs.

The ServerIron XL sends layer 3 (ping) health checks to each VPN and to the distant ServerIron. If a health check does not return through the same path, that link will be considered down and traffic will continue through only one VPN. For example, the ServerIron XL on site A will send health checks to the ServerIron XL on site B and C through both VPNs. If a health check does not return through VPN1, but it does from VPN2, the ServerIron XL will treat that link down. All traffic will only go through VPN2. The ServerIron XL will continue to send health checks through VPN1 to know when the link is operational. The amount of health checks, and the frequency can be manually specified.

The internal router should have one of the VPNs as its default gateway. When a packet arrives at the ServerIron XL, the packet will be load balanced accordingly. The ServerIron XL checks the source and destination IP address and runs it through a mathematical hash. Based on the hash, the packet is then forwarded to one of its VPNs. If that VPN were to fail, a new hash would be created before sending the traffic to its second VPN. When the first VPN is once again operational, the traffic is sent through the original VPN.

An internal router port must be defined which notifies the ServerIron of the location of the router. Even though the default gateway pointing to the router is specified, the ServerIron XL requires it. Each VPN is defined as a server. The ServerIron will treat each VPN as if they were firewalls. A group needs to be created to apply load balancing to. All VPNs that need to be load balanced must be a member of that group. By default, all ports are in fw-group 2.

Firewall paths define the path the health checks take to reach the distant end's ServerIron's IP address. If a VPN is not defined in the path, that link will always appear to be down. The ServerIron XL loadbalances traffic based on the user's configuration. To load balance all traffic, you have to enable load balancing of all Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) packets. BY simply configuring the ServerIron XL to load balance TCP and UDP, it is not necessary to specify applications or ports.

On the first port of the ServerIron XL, which has VPN1 connected, a manual entry has to be made that defines the MAC address of the VPN. Defining the static MAC entry prevents that entry from aging out of the configuration if the VPN becomes unavailable. It also keeps the ServerIron XL from having to ARP the VPN. Ensure that the MAC addresses entered correspond with the VPN connected to the port. Otherwise, communication with that VPN will fail. You can, but is not necessary, to add a static MAC entry for the router.

Zone Configuration

Multi-zone VPN LB allows you to configure the ServerIron XLs to forward packets based on the destination zone, or network. IT allows the ServerIron XL to have a complete topological view of the network. When you configure multi-zone VPN LB, you first identify a zone by configuring a standard Access Control Lists (ACLs). An ACL specifies the IP addresses (or address ranges) within the zone. It is NOT used to filter traffic. Unlike access lists applied to Cisco, they will not be applied to the interface. However, the access-list concept is the same.

When creating the access-list, place the busiest networks at the top of the list. The ServerIron XL supports wildcard bits for the subnet mask. If a mistake is done during the creation of the access-list, removing one entry will not remove the entire access-list. The number of the access-list must match the number specified during the creation of the zone.

When you configure the VPN group parameters, you add the zones and define them by associating the ACLs with them. Each zone consists of a zone number, an optional name, and a standard ACL, as previously mentioned, that specifies the IP addresses contained in the zone. When the ServerIron forwards a packet, it selects a path that goes through a VPN to a ServerIron that is in the zone that contains the destination IP address of the packet.

- Do not configure zone information on a ServerIron for the zone the ServerIron is in.
- On the remote ServerIron(s), configure zone definitions for the zone(s) in other networks.
- A ServerIron can only be a member of one zone.
- Do not configure zone 1. By default, this is the zone of last resort, if a ServerIron receives a packet destined for an unknown zone or network, it will send it to zone 1. According to Foundry, you can define zone 1 if you want to, but if you do, this zone contains only the IP address ranges you configure for the zone.⁵

Conclusion

Many companies are skeptical about using VPNs to secure their communication between multiple sites because of their low throughput. One way of eliminating this problem is by load balancing the traffic going through the VPNs. Currently, most VPNs support a cluster configuration in which multiple VPNs can be grouped together and viewed by the network as one. They share a virtual IP address that eliminates the single point of failure problem. However, only one VPN is used at any given time, therefore the bottleneck problem still exists. Most VPNs offer a high throughput when the security protocols used are weak. If a stronger standard is desired, it will result in performance degradation. In order to achieve the highest protection possible without having to deal with slow traffic, load balancing the VPNs can be the solution of choice. VPN load balancing, when used correctly, can solve most of the VPN's bandwidth and redundancy issues.

Bibliography

1. Alcatel. "Alcatel launches leading-edge windows 2000 solution for secure virtual private networks". About Alcatel's VPN Solutions. 29 Jan 2001. URL: <http://www.home.alcatel.com/vpr/> (13 Dec 2001).
2. DesAutels, Philip A. "SHA-1 Hash Algorithm - Version 1.0". Overview. 1 Oct 1997. URL: http://www.w3.org/TR/1998/REC-DSig-label/SHA1-1_0 (13 Dec 2001).
3. National Institute of Standards and Technology. "Federal Information Processing Standards Publication 46-3". Data Encryption Standard. 25 Oct 1999. URL: <http://csrc.nist.gov/publications/fips/fips46-3/fips46-3.pdf> (7 Dec 2001).
4. NETBSD. "IPsec FAQ". Transport mode and tunnel mode. 30 Aug 2001. URL: http://www.netbsd.org/Documentation/network/ipsec/#trans_tunnel (11 Dec 2001).
5. Foundry Networks. "Foundry ServerIron Firewall Load Balancing". Chapter 4 Configuring Multizone FWLB. 2001. URL: <http://www.foundrynet.com/services/documentation/siFWLB/ServerIronFWLBmultizone.html> (13 Dec 2001).

© SANS Institute 2000 - 2002, Author retains full rights.