



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>

Designing a Secure Local Area Network

Daniel Oxenhandler

GSEC – ver. 1.4b

Introduction

In order to design and build a well-secured network, many factors must be taken into consideration, such as the topology and placement of hosts within the network, the selection of hardware and software technologies, and the careful configuration of each component. My paper will be an examination of some of the issues in designing a secure Local Area Network (LAN) and some of the best practices suggested by security experts. I will discuss securing a LAN from the viewpoint of the network architect considering three main areas: the network topology which comprises the physical and logical design of the network; securing the routers and switches which connect segments and hosts to form the network; and, finally, some of the emerging and advanced techniques in network security will be examined.

Initial Assumptions and Challenges

My goal is to examine some of the security issues commonly found in the small to medium sized LAN set up for a business or other institution, and to identify some of the best practices from the perspective of the network designer. While no two networks are exactly alike, some of the typical challenges faced by the network designer include the following:

- Securing the network from Internet launched attacks
- Securing Internet facing web, DNS and mail servers
- Containing damage from compromised systems, and preventing internally launched attacks
- Securing sensitive and mission critical internal resources such financial records, customer databases, trade secrets, etc.
- Building a framework for administrators to securely manage the network
- Providing systems for logging and intrusion detection

Before beginning the design process, a security policy should be put in place, or updated to accurately reflect the goals of the company. Additionally, a realistic assessment of the risks faced, and identification of the resources (manpower, hardware, budget) that are available should be made. Once the organization's security policy and the available resources have been identified the design process can begin.

I have made the following assumptions for the sake of this discussion – we wish to secure a small to medium sized (under 500 hosts) TCP/IP based LAN which is connected to the Internet via broadband or other high speed connection. We

have a need for a reasonable amount of security because of mission critical records or proprietary information, but we are not guarding nuclear secrets or Fort Knox. Lastly, we will assume that we have adequate human resources and budget dollars to acquire and configure an optimum set of network technology. I will attempt to identify practices and technologies which can be tailored and applied appropriately to the individual site's needs.

Note: The sample configurations used in this paper are based on Cisco hardware because of their prevalence in the marketplace and my own familiarity with Cisco technology. The implementation details may differ from vendor to vendor, but most of the concepts discussed here will be transferable to hardware made by other manufacturers.

Topology and Architecture

A critical step in designing our network is defining the network topology. The topology is the physical and logical layout of the network. On the physical side, we will need to provide distribution to the offices or buildings where the users are located. We will need to provide connectivity to the servers which comprise our intranet, to the Internet, and possibly to other company locations or business partners, remote users connecting via telephone lines, etc. The logical topology must be considered as well. It is bound to some degree by the physical topology, but with technologies such as Virtual LANs (VLANs) and Virtual Private Networks (VPNs) there is considerable flexibility in designing the logical topology.

In laying out the logical topology we will need to consider our security policy, and decide what our trust model is. Which parts of the network are less trusted, and which are more? Which groups of devices and users should be logically grouped together, and which should be separated? Below is a graphic representing our initial network design:

© SANS Institute 2003




Figure 1: Basic Network Design

The basic design illustrates our connection to the Internet with a border router and firewall, and our public extranet servers which are connected to a third interface on the firewall. The firewall is one of four connections to a core router or, if higher performance is required, a layer 3 switch. The remaining connections to our core router are the floor or building switches which provide connectivity to the different departments and our intranet servers.

This topology demonstrates how devices with similar functions and security profiles are grouped together -- the public extranet servers, user workstations, and the intranet servers. By creating separate security zones, we will be able to enforce our security policy with the appropriate firewall rules and layer 3 access lists.¹

One element our basic design lacks is the infrastructure for managing our network. We will need one or more management workstations, tftp servers, and one or more syslog servers at a minimum. Other typical servers for the management network are a one-time password (e.g. RSA SecurID or Axent Defender) server, RADIUS server, etc. Because these servers will form the foundation of our network management and security, we will want to create a separate management VLAN which is isolated from the rest of the network by a

firewall or access lists. The only traffic that we will allow in to the management network is either from the managed devices or protected by encryption.

A design goal will be to keep management traffic off the production network, to eliminate the possibility that it could be intercepted in transit. Ideally we would configure each device with a physical port on the management VLAN. If this is not possible because of physical or other limitations, management should be encrypted via ssh or IPSEC. Below is a representation of the management network:




Figure 2: Management VLAN

Securing Routers and Switches

Now that the topology has been defined, let's take a look at building security into our network elements and configurations. Our design calls for segmenting the network into subnets based on function and, possibly, location. By implementing routing at the network core, our segments are isolated into individual broadcast domains. This improves performance and also improves security by preventing sniffing or arp based attacks between segments.

Within each subnet the hosts are connected to an Ethernet switch. A switch provides high performance by putting each host in its own collision domain, and enhances security by making sniffing and arp based attacks difficult. A hub is a less expensive alternative to a switch for layer 2 connectivity, though it is less desirable both from a performance and a security standpoint.

Layer 3 Design and Access Lists

Our layer 3 design is quite simple, with a central core router connecting the different production and management networks. Because we have mapped out our trust model and security policies, we can use access lists at layer 3 to implement our security policy. For traffic coming into a subnet, we will permit only appropriate incoming packets, based on the policy of that subnet. Similarly, we will filter outbound traffic to eliminate spoofing and minimize any malicious or illegitimate activities. Let's consider some example access lists based on the Cisco IOS command set.

Suppose we have a Windows 2000 file server and a web server on our server subnet. How do we configure our access list to permit the necessary traffic and deny everything else?

```
! Permit icmp echo to the server subnet (192.168.1.0/24)
! for troubleshooting
access-list 111 permit icmp any 192.168.1.0 0.0.0.255 echo
access-list 111 permit icmp any 192.168.1.0 0.0.0.255 echo-reply
! Permit Windows file sharing protocols to the Windows 2000
! server at 192.168.1.200
access-list 111 permit tcp any host 192.168.1.200 eq 135
access-list 111 permit tcp any host 192.168.1.200 eq 139
access-list 111 permit tcp any host 192.168.1.200 eq 445
access-list 111 permit udp any host 192.168.1.200 eq 137
access-list 111 permit udp any host 192.168.1.200 eq 138
access-list 111 permit udp any host 192.168.1.200 eq 445
! Permit http access to the web server at 192.168.1.201
access-list 111 permit tcp any host 192.168.1.201 eq 80
! Deny any other traffic
access-list 111 deny ip any any log
```

The above commands illustrate the concept of our layer 3 design, and would need to be expanded and modified in a production environment. Let's now consider a workgroup subnet populated with desktops but no servers. Since we don't expect servers to be placed here, inbound tcp traffic is limited:

```
! Permit icmp echo to the workgroup subnet (192.168.2.0/24)
! for troubleshooting
access-list 121 permit icmp any 192.168.2.0 0.0.0.255 echo
access-list 121 permit icmp any 192.168.2.0 0.0.0.255 echo-reply
! Permit established tcp connections only
access-list 121 permit tcp any 192.168.2.0 0.0.0.255 established
! Permit inbound udp traffic to support NetMeeting
access-list 121 permit udp any 192.168.2.0 0.0.0.255 gt 1023
! Deny any other traffic
```

```
access-list 121 deny ip any any log
```

Finally, we will want to filter traffic leaving each subnet to prevent spoofing. The presence of incorrect source addresses could indicate either a misconfigured machine, or one which was compromised and attempting to launch a DDOS or similar attack. Here's how outbound filters would be defined for the workgroup subnet:

```
! Permit outbound traffic from workgroup subnet
! (192.168.2.0/24) with a legitimate source address;
! Deny all other traffic
access-list 122 permit ip 192.168.2.0 0.0.0.255 any
access-list 122 deny ip any any log
```

Securing Layer 3

We have illustrated above how our layer 3 design and access lists are used to implement our security policies. We also want to take steps to ensure that the routers themselves are secured against attacks. There are many excellent templates for hardening Cisco routers against attacks such as the National Security Agency's or Cisco's.^{2,3} I would like to point out a couple key strategies that are relevant to our discussion.

Secure management of the routers is enforced by several mechanisms. First is the management VLAN, which ensures that the management traffic does not traverse the production network. Access lists should be configured on the management ports to block illegitimate connections. Out Of Band (OOB) communication, such as via a terminal server, is another excellent means of securing management traffic. We will use strong authentication provided by a one-time password server, such as RSA Security's ACE server. Encrypted communication protocols such as ssh should be used if in band (over the production network) communication is necessary. Logging to the syslog servers located on the management network will meet our auditing requirements.^{4,5}

The following excerpt from a Cisco 7500 router configuration file demonstrates some of these concepts:

```
! Configure logging to syslog server (192.168.7.88)
logging trap debugging
logging facility local7
logging 192.168.7.88
! Configure aaa authentication to the radius server
! on the management VLAN
aaa new-model
! create a local user account in case radius
! is unavailable
user freduser password [password]
aaa authentication login default group radius local
radius-server host 192.168.7.77 auth-port 1645
radius-server timeout 5
radius-server key [RADIUS shared secret]
```

```
! Configure ssh server
crypto key generate rsa
ip ssh time-out 60
ip ssh authentication-retries 2
! Access list to ensure management traffic is sourced from
! the management station (192.168.7.23)
access-list 7 permit host 192.168.7.23
access-list 7 deny ip any any
! Configure management port
line vty 0 4
access-class 7 in
transport input ssh
```

Layer 2 Design

In previous sections, we have described how security can be implemented in the layer 3 design via access lists and hardening the routers themselves. We must now address threats that exist at layer 2 and continue to enforce our security policy in the layer 2 design.

One question we will want to consider is how to maximize the security of the switch ports themselves. If an attacker controlled a host on one of our VLANs, could she jump to another VLAN and gain access to a more sensitive VLAN? What about the possibilities of a misconfiguration providing undesired access to an intruder? To achieve the highest level of security we would configure only one VLAN per switch.⁶ This would minimize the chance of an attacker jumping VLANs and reduce the chance of misconfiguration. If we can provide this kind of isolation of one VLAN per switch, it is the most secure, and highly recommended for likely attack points such as our Internet facing server segment.

A July, 2000 study by David Taylor found that it was possible to jump across VLANs on a Cisco switch by injecting specially crafted frames in a default configuration.⁷ A more recent study commissioned for Cisco by researchers at @Stake proclaimed that there was minimal risk in using VLANs when configured according to best practices.⁸ Depending on her budget and confidence in her configurations, the designer may decide that it is okay to combine multiple VLANs on a single switch. The designer will have to carefully weigh the costs and risks, and make a decision appropriate to her environment.

Since the switch ports are the gateway into our network, we will want to implement physical security when possible, by controlling access to switch ports, and disabling unused ports. As most busy network admins may not be able to monitor every unused port, there are many other techniques that can be used to enhance security. One technique is to require the users to authenticate via RADIUS or LDAP before they are given access to any resources. This technology is implemented in Cisco's *User Registration Tool* (URT) or competing vendor Extreme Networks *network login* feature.^{9,10} Cisco's URT allows users to be assigned to different VLANs depending on the credentials supplied.

However, implementation of this feature involves a complex infrastructure, and is best suited to very large enterprises.

Many other strategies for securing access to switch ports are available. Limiting the MAC addresses that are permitted to communicate on the ports is key to layer 2 security. A flood of MAC addresses, or even a single new MAC address could indicate an intruder, or ARP spoofing activities such as the *dsniff* utility.¹¹ Many switches can be configured with static or secure MAC assignments. Creating a static MAC assignment ensures that frames for the designated ethernet address are always forwarded to the specified port, and it can prevent ARP spoofing attacks. To set a static port on a Cisco Catalyst switch (CatOS 6.3) the following statement is used:

```
set cam permanent aa-bb-cc-11-22-33 6/1
```

The preceding command will ensure that frames for the specified MAC will always be forwarded to the specified port.¹² Static MAC assignments are especially advisable for critical hosts like gateways and firewalls.

Another good idea is to limit the number of MAC addresses that can appear on each port, either to one or an appropriate small number, or configure a timeout that prevents a new MAC from appearing until a certain time period elapses.¹³ These features can be configured with the *set port security* statement on a Cisco Catalyst switch. The following statements show some of the options available with this command:

Limit the number of permissible MAC addresses to 1

```
set port security 6/1 enable maximum 1
```

Limit the permissible MAC on port 6/1 to aa-bb-cc-11-22-33

```
set port security 6/1 enable aa-bb-cc-11-22-33
```

Limit the permissible MAC addresses to 2, and the aging to 30 minutes

```
set port security 6/1 enable age 30 maximum 2
```

Limit the permissible MAC addresses to 1, aging to 30 minutes, and the violation action to restrict packets from insecure hosts (default is shutdown the port)

```
set port security 6/1 enable age 30 maximum 1 violation restrict
```

Securing Layer 2

We have surveyed some of the key strategies for securing our network from violations of our layer 2 policy. But we must take additional steps to secure the switches themselves from attacks, and from attacks against layer 2 protocols such as Spanning-Tree Protocol (STP), which could result in denial of service or availability situations.

Locking down the security on the layer 2 devices will follow the principles used to lock down the routers, such as disabling insecure default configurations, securing the management channel and implementing strong authentication via one-time passwords. Once we have hardened the switch itself, we will want to secure our infrastructure against attacks on the underlying layer 2 protocols.

Spanning-Tree Protocol (STP) is used by switches and bridges to establish their MAC address forwarding tables, and establish a tree-like topology which forwards frames via the fastest path and eliminates loops. Bridge Port Data Units (BPDUs) are exchanged by switches to share information about the topology.

STP is designed to dynamically adjust to changes in the topology, but it is vulnerable if random hosts start transmitting BPDUs and affecting the spanning tree. This could happen if an unauthorized switch was attached to one of our ports, or a bridging protocol was enabled on a Linux host, for example. If our switch supports it, we can prevent random hosts from either forwarding BPDUs or affecting the spanning tree. Cisco Catalyst switches provide two features which address this problem, *STP Root Guard*, and *STP Portfast BPDU Guard*.

For optimum performance, we will want the root bridge of the spanning tree to be located near the core of the network on the highest bandwidth links. The STP root guard feature allows us to enforce the STP topology, and prevent the root bridge from appearing on an edge segment, or on a lower bandwidth connection. Root guard is enabled on ports where we don't want the root to appear. If superior BPDUs are received from a port with root guard enabled, the port will change from forwarding to listening state until the superior BPDU announcements are stopped.¹⁴ Root guard is enabled as follows:

```
set spantree guard root 5/1
```

The root guard feature allows us to control ports where we do expect to receive STP announcements, and is configured on a per-port basis. This makes it suited to ports which are connected to other switches. On ports designated for host access, we will use the *spanning tree portfast BPDU guard* feature.¹⁵

In a normal STP configuration, when a host starts transmitting on a port, a spanning tree calculation is performed which takes 30-50 seconds before the port enters the forwarding state and begins transmitting frames. The *spanning tree portfast* command is typically configured on ports where end stations are attached, and allows the port to immediately transition to the forwarding state, without the delay caused by the STP calculation. However, a port configured with the portfast feature still participates in STP, and there is the possibility that a device communicating on that port will affect the spanning tree topology and the placement of the root bridge.

The spanning tree portfast BPDU guard feature addresses this issue, by disabling the port if a BPDU is received on that port. This feature is enabled for all portfast enabled ports as follows:

```
set spantree portfast bpdu-guard enable
```

Advanced Technologies

So far, we have discussed the physical and logical topology of our LAN, and techniques for building security into the layer 3 and layer 2 design. There is still much that we can do to make our design more secure. In this final section on Advanced Technologies, we will look at a few more technologies used by security professionals to detect and deter crackers.

Intrusion Detection Systems

In this section we will briefly consider Network Intrusion Detection Systems (NIDS) and how they can be used in our LAN to detect undesirable activity. Many experts would include IDS as part of the essential elements of securing any network. Network IDS can alert the system administrator to attacks on the network in real time by inspecting the traffic on the wire, and generating alerts if suspicious activities are identified. NIDS can be a regular computer running IDS software, such as the freeware Snort, an appliance type device running proprietary software, or even a specialized card built in to a switch or other network element as Cisco has recently introduced. Host based intrusion detection, such as free or commercial versions of Tripwire, or various kinds of proactive log monitoring software, are also highly recommended, but outside the scope of this paper.

Once we have selected a NIDS for use in our network, we will need to place the sensors logically within the topology. Unless we have lots of resources for maintaining our NIDS, and analyzing and responding to alerts, we will want to limit ourselves to a few well placed sensors. Because we have a switched infrastructure, we will need to connect the NIDS sensors to a specially configured monitoring port where all the traffic from a VLAN is mirrored. In a high bandwidth environment there will be physical limits to the IDS system that will need to be considered as well – a standard PC running IDS software will not be able to keep up with a highly loaded gigabit ethernet VLAN, for example. An alternate and more complex solution is to use network taps, a hardware component specialized for monitoring network connections.¹⁶ Below is a figure which shows how the IDS sensors could be placed in our network:




Figure 3: Location of IDS Sensors

We place a sensor on our Internet facing segment, because the public servers are a visible target to attackers. Another sensor is placed behind the firewall, to monitor traffic between the Internet and our internal LAN. If we had a remote access segment, for instance where a dial-up server or VPN concentrator terminated, this would also be a good place for an IDS sensor. Finally, we will want to locate a sensor on the more sensitive subnets within our network, the intranet server subnet and the management subnet. An attack on either of these segments could have very serious consequences. We may also wish to place a sensor outside our firewall to monitor what kinds of attacks are being launched against us, but which are screened out by the firewall. This sensor, if used, will generate the most data and false positives, so the sensitivity should be adjusted accordingly.

Private VLANs and VLAN ACLs

In previous sections we discussed several strategies for securing our layer 2 and layer 3 design. However, as the sophistication of our defenses improves, it is inevitable that skilled attackers will find new techniques to defeat our defenses in the ever-changing "Information Warfare" environment. So the security professional must keep aware of emerging techniques to defend her network. Two features offered by Cisco on their high-end switches are worthy of consideration in this regard, Private VLANs and VLAN ACLs.¹⁷

Private VLANs allow the designer to enforce a security policy within a subnet. For instance if host A and host B are on the same subnet, typically there is nothing that prevents them from communicating with each other. However, there may be situations where this behavior is not desirable, such as on our Internet facing segment. If a cracker was able to gain entry to one of our public servers, he would logically launch attacks against other hosts on the public segment. Private VLANs provide a means to prevent hosts on the same subnet from communicating with each other, while permitting required communication to their router and hosts on other networks.

Private VLANs are established by defining a primary VLAN and one or more secondary VLANs on a segment. Hosts defined on an "isolated" secondary VLAN will only be permitted to communicate to their gateway, and will be prevented from talking to other hosts on the same primary VLAN. Note that host based firewalls could be used to achieve similar results, and are available built in to many operating systems, or from numerous software vendors.

VLAN ACLs (VACLs) are another security enhancement offered on Cisco's high-end switches. VLAN ACLs can further enhance the security afforded by PVLANS by preventing undesired traffic sourced from the VLAN. VACLs also prevents a compromised or misconfigured host from using the gateway to communicate to other hosts on the same PVLAN. Because the ACL processing takes place in the switch ASICs, Cisco's claim is that it can happen at wire speed, thus invoking no performance penalty in high throughput environments. Security is also enhanced because traffic is denied at layer 2, before it even passes to the router for processing.

Micro VLANs

Micro VLANs or Routing to the Desktop (R2D) is another emerging strategy that has been proposed.¹⁸ In this design, each host is placed in its own routed VLAN on a layer 3 switch. This design eliminates the vulnerabilities in spanning tree protocol, arp spoofing, and attacks on Hot Standby Routing Protocol (HSRP). If each port was assigned its own VLAN with a 30-bit subnet mask, there would only be one valid host IP address that could appear there. This would reduce the risks associated with IP address spoofing, or the introduction of rogue machines on the network. While the concept of Micro VLANs is radical, it is certainly

feasible with hardware available today. In specialized environments where a high degree of isolation and security are needed, Micro VLANs may fit the bill.

IPSEC

A final technique that should be considered is implementing security at the network level. Strong encryption and authentication implemented at the network level would prevent all but the most determined attacker from compromising our hosts, even if he were able to penetrate our perimeter defenses. IP Security (IPSEC) is an enhancement to the IP protocol documented in various RFCs by the IETF. IPSEC ensures that every packet transmitted on the LAN is encrypted with strong encryption algorithms.

While IPSEC has been criticized for its complexity, it is emerging as the standard for network level encryption.¹⁹ IPSEC is the underlying protocol used by many of the VPN solutions currently on the market. And it is finding use in high security environments, or in special applications such as management of network devices over insecure networks. Within the last few years IPSEC has become widely supported in popular operating systems and broader adoption seems inevitable. Microsoft has included IPSEC in its ubiquitous operating systems, along with an easy to use configuration wizard.

Conclusion

This paper has examined several strategies for designing a secure Local Area Network. We have identified the need to define a security policy, and balance the organization's security needs with the available resources. Next, we considered a basic topology that allows for the grouping of hosts by function, and implementing security within the layer 3 design. We looked at how to implement the layer 2 network design securely, implementing layer 2 security features and minimizing threats at layer 2. And finally, we rounded out our discussion by discussing additional steps we can take to secure the LAN, such as network intrusion detection systems, private VLANs and IPSEC. The network architect must exercise careful planning, and attention to detail to maximize the security of the network while meeting the communication needs of the organization.

List of References

¹ Convery, S., Trudel, B. "SAFE: A Security Blueprint for Enterprise Networks." 2000. URL: http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/safe_wp.pdf (11 Dec. 2002).

² Antoine, Vanessa, et al. "Router Security Configuration Guide." Version 1.1. 27 September 2002. URL: <http://www.nsa.gov/snac/cisco/guides/cis-2.pdf> (11 Dec. 2002).

³ "Improving Security on Cisco Routers." 29 Oct. 2002. URL: <http://www.cisco.com/warp/public/707/21.html> (15 Dec. 2002).

⁴ Convery, S., Trudel, B.

⁵ "Cisco ISP Essentials: Essential IOS Features Every ISP Should Consider." Version 2.9. 6 June 2001. URL: <http://www.cisco.com/public/cons/isp/documents/IOSEssentialsPDF.zip> (11 Dec. 2002).

⁶ Stephens, Brian. "Architecting Secure Network Topologies." 10 Dec. 2001. URL: http://dcb.sun.com/practices/howtos/network_topologies.jsp (15 Dec. 2002).

⁷ Taylor, David. "Are there Vulnerabilities in VLAN Implementations?" Intrusion Detection FAQ. 12 July 2000. URL: <http://www.sans.org/newlook/resources/IDFAQ/vlan.htm> (15 Dec. 2002).

⁸ Pollino, D., Schiffman, M. "Secure Use of VLANs: An @stake Security Assessment." Aug. 2002. URL: http://www.cisco.com/warp/public/cc/pd/si/casi/ca6000/tech/stake_wp.pdf (15 Dec. 2002).

⁹ "Data Sheet: Cisco Secure User Registration Tool Version 2.5." 29 Oct. 2002. URL: http://www.cisco.com/warp/public/cc/pd/wr2k/urto/prodlit/urt_ds.htm (15 Dec. 2002).

¹⁰ "HIPAA Compliance: An Extreme Approach." 2002. URL: http://www.extremenetworks.com/libraries/whitepapers/technology/HIPAA_WP.asp (15 Dec. 2002).

¹¹ "dsniff" URL: <http://naughty.monkey.org/~dugsong/dsniff/> (15 Dec. 2002).

¹² "Catalyst 5000 Family Command Reference (6.3)." 21 Mar. 2002. URL: http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/rel_6_3/cmd_ref/sets_sete.htm#1062845 (15 Dec. 2002).

¹³ Gill, Stephen. "Catalyst Secure Template." 1 Nov. 2002. Version 1.21. 14 Nov. 2002. URL: <http://www.qorbit.net/documents/catalyst-secure-template.htm> (15 Dec. 2002).

¹⁴ "Spanning-Tree Protocol Root Guard Enhancement." 29 Oct. 2002. URL: <http://www.cisco.com/warp/public/473/74.html> (15 Dec. 2002).

¹⁵ "Spanning Tree Portfast BPDU Guard Enhancement." 21 Nov. 2002. URL: http://www.cisco.com/en/US/tech/tk389/tk689/technologies_tech_note09186a008009482f.shtml (15 Dec. 2002).

¹⁶ Laing, Brian. "How To Guide: Intrusion Detection Systems." 2000. URL: <http://www.snort.org/docs/iss-placement.pdf> (15 Dec. 2002).

¹⁷ "Securing Networks with Private VLANs and VLAN Access Control Lists." 29 Oct. 2002. URL: <http://www.cisco.com/warp/public/473/90.shtml> (15 Dec. 2002).

¹⁸ Dugan, Stephen. "Putting 2 and 2 Together: Designing Security into Your Network Infrastructure." 2002. URL: <http://www.blackhat.com/presentations/bh-usa-02/bh-us-02-dugan-layer.ppt> (15 Dec. 2002).

¹⁹ Ferguson, F., Schneier, B. "A Cryptographic Evaluation of Ipsec." Feb. 1999. URL: <http://www.counterpane.com/ipsec.pdf> (15 Dec. 2002).

© SANS Institute 2003, Author retains full rights.