Network Security Architectures

Expert guidance on designing secure networks

ciscopress.com

Sean Convery, CCIE® No. 4232

Network Security Architectures

Sean Convery, CCIE No. 4232

Cisco Press

Cisco Press
800 East 96th Street
Indianapolis, IN 46240 USA

In the context of this discussion, *data center* refers to any location where centralized network resources are stored. This could include traditional data centers, wiring closets, coat closets, or someone's desk. It all depends on the size of the facility and the way it is organized.

TIP Some ultrasecure data centers utilize sets of cameras, key card access, biometrics, and "man-traps" to catch anyone illegally trying to gain access to the room.

Separate Identity Mechanisms for Insecure Locations

Although identity design considerations are discussed in more detail in Chapter 9, "Identity Design Considerations," from a physical security perspective, it is important to ensure that passwords in physically insecure locations are not the same as those used in secure locations.

Often an organization will utilize common authentication mechanisms for the various systems that must access network resources. For example, SNMP community strings or Telnet/ SSH passwords might be set the same on all devices. From a pure security perspective, it is preferable to use two-factor authentication, when available, for each user who accesses the network device. Although this might be possible for users, it is often impossible for software management systems, which need to run scripts to make changes on several machines at once. For optimal security, different passwords should be used on each device, but this is often operationally impossible for large networks.

Therefore, at a minimum, organize your common passwords so that they are never used on systems in physically insecure locations. For example, assume you have 3 main locations (with data centers) to your organization and 10 remote sites (considered insecure). In this case, only use your shared passwords on the main sites and ensure that the passwords for each of the remote systems are unique per site at a minimum and per device ideally. As the number of insecure locations increases into the hundreds or thousands, this becomes impossible; refer to the "Business Needs" section of Chapter 2, "Security Policy and Operations Life Cycle," for guidance on calculating the costs and benefits of this and any other difficult security measure. (People generally don't compute cost/benefit on easy and cheap security measures.)

Prevent Password Recovery Mechanisms in Insecure Locations

Some devices have controls to prevent the recovery of passwords in the event that an attacker has physical access to your system. For example, on some newer Cisco routers and switches, the command is as follows:

```
Router(config)# no service password-recovery
```

When this command is entered on a router or a switch, interrupting the boot process only allows the user to reset the system to its factory default configuration. Without this command, the attacker could clear the password and have access to the original configuration. This is

important because the original configuration might contain common passwords or community strings that would allow the attacker to go after other systems.

This would be particularly useful in insecure branch offices or other locations where the physical security of a network device cannot be assured.

Be Aware of Cable Plant Issues

In today's networks, there are two primary cable types: unshielded twisted pair (UTP) category 5 (or higher) and fiber optic. The risk of an attacker accessing your physical cabling is important to consider because that level of access often can bypass other security controls and provide the attacker with easy access to information (provided encryption is not used). UTP cable is very easy to tap, but it was thought years ago that fiber was immune to cable taps. We now know that this is not the case. The National Security Association (NSA) is rumored to have already tapped intercontinental network links by splicing into the cable; read about it at the following URL: http://zdnet.com.com/2100-11-529826.html.

It is also theorized that fiber cable could be bent far enough so that some light would escape if the outer layer of the cable is removed. With the right types of equipment, this information could then be read.

Additionally, if an attacker gains physical access to a wiring closet or the fiber cable as it runs in a cable tray above a drop ceiling, tapping the cable by installing couplers is another possibility.

All this being said, fiber is more secure than copper because the means to tap the signal are more expensive, difficult to execute, and often require interrupting the original flow of data to install. On the other hand, the means to tap a UTP signal can easily be purchased off of the Internet.

Be Aware of Electromagnetic Radiation

In 1985, the concerns of the paranoid among the security community were confirmed. Wim van Eck released a paper confirming that a well-resourced attacker can read the output of a cathode-ray tube (CRT) computer monitor by measuring the electromagnetic radiation (EMR) produced by the device. This isn't particularly easy to do, but it is by no means impossible. Wim's paper can be found here:

http://www.shmoo.com/tempest/emr.pdf

This form of attack is now commonly called *van Eck phreaking*. Additionally, in 2002, Markus Kuhn at the University of Cambridge published a similar method of reading data off of a CRT, this time by measuring the changes in the amount of light in a room. His paper can be found here:

http://www.cl.cam.ac.uk/~mgk25/ieee02-optical.pdf

And an easy-to-read FAQ on the topic can be found here:

http://www.cl.cam.ac.uk/~mgk25/emsec/optical-faq.html

A simple way to mitigate van Eck phreaking might just be to change the type of font you are using. Ross Anderson and Markus Kuhn did some excellent research on the topic:

http://www.cl.cam.ac.uk/~mgk25/ih98-tempest.pdf

I am certainly not recommending that all systems must address these sorts of security considerations, but it is good to know that such attacks are possible.

Be Aware of Physical PC Security Threats

Oftentimes, inexperienced network designers begin with an unacknowledged assumption that *all* the sensitive data within an organization is contained on servers. In reality, there is sensitive information about my company sitting on the laptop I am using to write this book, as well as on the servers. Like most employees at my company, server resources are used when necessary, but often interesting information is stored locally.

Several physical security issues manifest when you operate under the preceding assumption:

- The first is that portable computer theft is a big problem, not just in the cost of replacing the computer but in the proprietary information that is stored on it. The best protection against having a lost portable computer turn into lost trade secrets is some type of file system encryption. (Some are built into modern OSs.) Chapter 4 has more details on such systems.
- The second is that by compromising the data coming into and out of a PC, you can learn passwords, sensitive data, and so on. An attacker can achieve this through network sniffing, EMR emissions (discussed in the previous section), remote control software (Back Orifice 2000), or novel devices that attach between the keyboard and the PC and record to flash memory every key typed. For more information see this URL:

  http://www.thinkgeek.com/stuff/gadgets/5a05.shtml

Layer 2 Security Considerations

As you learned in Chapter 3, "Secure Networking Threats," certain attacks run at Layer 2 (L2) of the OSI model. Oftentimes, your posture toward L2 attacks depends on the physical security of the location and the amount of trust you have in users, as defined by your security policy. This section discusses some common design considerations for L2 protocols. The discussion is focused on Ethernet, but most of these issues apply to wireless networks as well.

L2 Control Protocols

Control protocols are usually at the core of any L2 security issue. This section discusses design considerations around L2 control protocol usage. Basic understanding of these protocols is assumed. There are two main topics in this section: the first covers industry-standard protocol considerations; the second covers Cisco-specific protocols.

General Protocol Considerations

This section covers the standard protocols 802.1q, Spanning-Tree Protocol (STP), and briefly mentions 802.1x.

802.1q

The 802.1q standard specifies a standard mechanism for Ethernet switches to exchange virtual LAN (VLAN) information. It adds a 4-byte tag after the source and destination Media Access Control (MAC) addresses. The first 2 bytes act as an Ethernet tag protocol identifier. The second 2 bytes contain all the interesting information. Twelve bits are used as a VLAN identifier (yielding 4096 choices), and 3 bits are used as a priority identifier (in the 802.1p standard). The addition of 4 bytes to the Ethernet packet increases the maximum size of an Ethernet frame from 1518 bytes to 1522 bytes.

When designing a network to take advantage of 802.1q tagging, there are a few security concerns that must be addressed:

- 802.1q has had several implementation flaws in various vendors' equipment over the years. Details of an old Cisco vulnerability can be found here: http://www.sans.org/resources/idfaq/vlan.php. Many of these problems have been fixed, and vendors are beginning to pay more attention to security, particularly as VLANs play a greater role in any network design.
- When using VLANs, the potential for human error increases because the operator must keep track of "virtual" LANs that might not have distinct cable plants associated with them. This can get particularly nasty when you try to remember which VLAN number is the outside of your firewall as opposed to the inside. Good management tools can mitigate the impact of this concern.
- Some attacks that use 802.1q as an attack method are detailed in a later section of this chapter titled "VLAN Hopping Considerations."

STP

Spanning-Tree Protocol (STP) is a L2 loop avoidance mechanism. Without STP, redundant L2 links would cause large forwarding loops and massive performance problems. From a security standpoint, STP has a few design characteristics of interest.

First, STP has no provisions for authentication of the bridge protocol data units (BPDUs) that are sent from switches and bridges as they exchange STP information. These BPDUs could easily be sent from an unauthorized device that could have any number of undesirable effects.

To start with, if the attacker can cause a failure of a link in the forwarding state, it generally takes 30 to 45 seconds for STP to deal with the failure and reconverge the topology. Some switches now include features to deal with this problem. On Cisco devices, the features are called port fast and uplink fast.

Second, for there to be some "authority" in the STP network, the participating switches elect a *root* bridge. It is from this bridge that the *loop-free* topology is built. The method for determining the root bridge is generally through STP configuration messages, which indicate the bridge priority of a given switch. The lowest number becomes the root bridge. If an attacker is able to send out BPDUs from his station, he can send out a configuration message with a bridge priority of zero. This will likely make his system the root bridge and will often change which links are active on a given network (since the topology is redetermined from the perspective of the new root bridge). No special tools are needed to do this; some UNIX implementations come with Ethernet bridging utilities that allow them to configure their system as a bridge with full participation in the STP process. As an example, consider the following topology in Figure 6-1.

Figure 6-1 *Starting Topology*

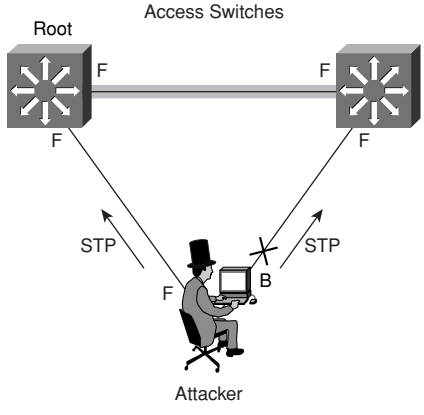

In the figure, you can see that the attacker has established two links to two different L2 switches. F denotes a link that is forwarding; B is a link that is blocked because of STP. This could easily be done by walking a long cable to another jack in a building or by using a WLAN network (if it was poorly designed). From here, you can see that one of the attacker's links is in the blocking state. This is exactly what STP should do to prevent loops. However, the attacker then sends BPDUs advertising himself as bridge priority zero. This causes STP to reconverge and the attacker to become the root bridge. A topology that looks like the one in Figure 6-2 results.

Figure 6-2 *Resulting Topology*

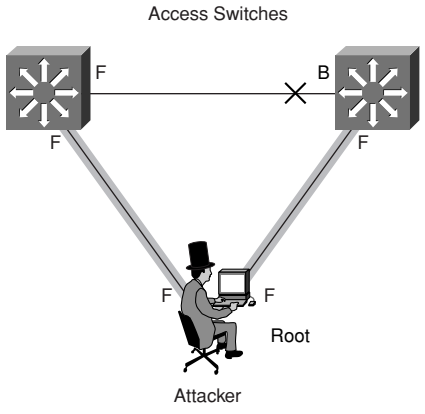

Because the topology is built from the perspective of the attacker, you can see that all traffic that must pass between the switches flows through the attacker's PC. This allows an attacker any number of options, as outlined in Chapter 3. The most obvious are sniffing traffic, acting as a man-in-the-middle, or creating a denial of service (DoS) condition on the network. The DoS condition is achieved because the attacker can make his links much slower than the links between the two access switches, which could very likely be connected by gigabit Ethernet.

NOTE You might ask, "Doesn't STP take into account bandwidth speed when determining the topology?" It does but always from the perspective of the root bridge. While testing in the lab, I was able to take a full-duplex gigabit link between two access switches and reduce it to a half-duplex 10 megabit (Mb) connection between those access switches and the attacking PC. This is never good for a production network.

Fortunately, mitigating this attack is fairly straightforward. First, some advocate disabling STP in all cases in which you don't have network loops. Although this sounds like a good idea, the attacker could instead introduce a loop into your network as a means of attack. A better option is to filter which ports are allowed to participate in the STP process. Some switches offer the ability to do this today. On Cisco devices, the two principal options are BPDU Guard and Root Guard.

BPDU Guard BPDU Guard can be globally enabled on some Cisco switches and is in effect on any port configured with the port fast option. Port fast ports are generally user ports. What BPDU Guard does is disable any port fast port that receives a BPDU message. Because these are user ports, there should be no reason for BPDU messages to be sent to them. The syntax is as follows: