Fundamental Principles of Network Security

By Christopher Leidigh

White Paper #101

Executive Summary

Security incidents are rising at an alarming rate every year. As the complexity of the threats increases, so do the security measures required to protect networks. Data center operators, network administrators, and other data center professionals need to comprehend the basics of security in order to safely deploy and manage networks today. This paper covers the fundamentals of secure networking systems, including firewalls, network topology and secure protocols. Best practices are also given that introduce the reader to some of the more critical aspects of securing a network.

©2005 American Power Conversion. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted, or stored in any retrieval system of any nature, without the written permission of the copyright owner. *www.apc.com* Rev 2005-0

Introduction

Securing the modern business network and IT infrastructure demands an end-to-end approach and a firm grasp of vulnerabilities and associated protective measures. While such knowledge cannot thwart all attempts at network incursion or system attack, it can empower network engineers to eliminate certain general problems, greatly reduce potential damages, and quickly detect breaches. With the ever-increasing number and complexity of attacks, vigilant approaches to security in both large and small enterprises are a must. **Figure 1** illustrates the steep rise in security incidents occurring each year, as reported to the CERT® Coordination Center (a center of Internet security expertise).

Figure 1 *– Security incidents by year – CERT.ORG*

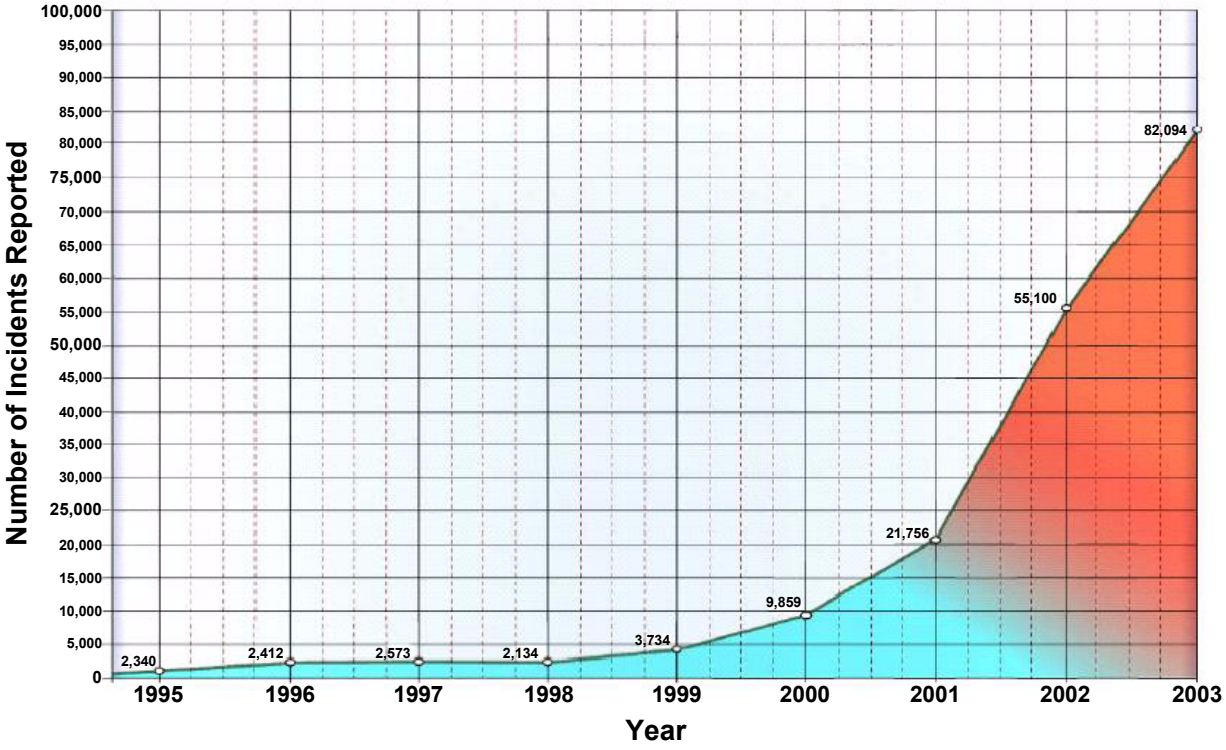

© 1998-2003 by Carnegie Mellon University

This paper presents security fundamentals as well as some best practices regarding the network, computer hosts and infrastructure network elements. Since there is no such thing as "the only way to approach security", it is left up to the reader / implementer to best judge what measures are appropriate.

The people problem

People are truly the weakest link in any security schema. Most people are not careful about keeping secrets such as passwords and access codes that form the basis for most secure systems. All security systems rely on a set of measures employed to control access, verify identity and protect disclosure of sensitive information. These measures usually involve one or more "secrets". Should a secret be revealed or stolen

©2005 American Power Conversion. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted, or stored in any retrieval system of any nature, without the written permission of the copyright owner. *www.apc.com* Rev 2005-0

then the systems that are protected by these secrets can be compromised. It may seem like a terribly obvious statement, but most systems are compromised in very basic ways. Leaving a Post-It note with a system password stuck to the side of a computer monitor may seem foolish, but many people in fact do such things. Another example, which is only slightly less obvious, is the tendency to leave factory default passwords in certain network devices. One such device might be a network management interface to a UPS. UPS systems, whether small in capacity or large enough to power 100 servers, are often overlooked in a security scheme. If such devices are left with default usernames and passwords, it could just be a matter of time before someone gains access knowing nothing more than the device type and its published default credentials. Imagine a server bank with rock solid security protocols on each web and mail server crashed by a simple power cycle on an unprotected UPS!

Security, the big picture

A secure “enterprise”, big or small, should have an approach to security that is comprehensive and end-to-end if it is to be effective. Most organizations do not have such policies and practices in place. There are some good reasons for this; security clearly comes at a cost. This cost can be measured not just in dollars, but also in complexity, time and efficiency. To make things secure, it is necessary to spend money, perform more procedures, and wait for these procedures to complete (or perhaps involve someone else).

The reality is that true security programs are difficult to achieve. It is usually necessary to choose a schema that has a certain amount of “cost” and an understood amount of security coverage. (This is almost always less than “comprehensive and end to end”.) The point here is to make educated decisions for each aspect of an overall system and to consciously employ more or less in a calculated fashion. If one knows the areas that are less protected, one can at least monitor such areas to determine problems or breaches.

Security Basics

Knowing the network

It is not possible to protect anything unless one clearly understands WHAT one wants to protect. Organizations of any size should have a set of documented resources, assets and systems. Each of these elements should have a relative value assigned in some manner as to their importance to the organization. Examples of things that should be considered are servers, workstations, storage systems, routers, switches, hubs, network and Telco links, and any other network elements such as printers, UPS systems and HVAC systems. Other important aspects of this task include documenting equipment location and any notes on dependencies. For instance most computers will rely on power backup systems such as UPSs which themselves may be part of the network if they are managed. Environmental equipment such as HVAC units and air purifiers may also be present.

©2005 American Power Conversion. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted, or stored in any retrieval system of any nature, without the written permission of the copyright owner. *www.apc.com* Rev 2005-0

Understanding different threats

The next step is to identify the potential "threats" to each of these elements as shown in **Table 1**. Threats can come from both internal and external sources. They may be human based, automated or even non-intentional natural phenomenon. The latter might more appropriately be categorized under system health threats as opposed to security threats, but one issue can lead to the other. One example is a power outage to a burglar alarm. The power outage could be intentional or through some natural event such as a lightning strike. In either case security is diminished.

Table 1 *– Summary of various threats with consequences*

Threat	Internal \ External	Threat Consequences
e-mail with virus	External origination, internal use	Could infect system reading email and subsequently spread throughout entire organization
Network virus	External	Could enter through unprotected ports, compromise whole network
Web based virus	Internal browsing to external site	Could cause compromise on system doing browsing and subsequently affect other internal systems
Web server attack	External to web servers	If web server is compromised hacker could gain access to other systems internal to network
Denial of service attack	External	External services such as web, email and ftp could become unusable If router is attacked, whole network could go down
Network User Attack (internal employee)	Internal to anywhere	Traditional border firewalls do nothing for this attack. Internal segmentation firewalls can help contain damage.

Physical security, protection on the inside

Most experts would agree that all security starts with physical security. Controlling physical access to machines and network attach points is perhaps more critical than any other aspect of security. Any type of physical access to an internal site creates a major exposure of the site. Secure files, passwords, certificates and all sorts of other data can usually be obtained if physical access is possible. Fortunately there are all sorts of access control devices and secure cabinets that can help with this problem. For more information on physical security of data centers and network rooms see APC White Paper #82, "Physical Security in Mission Critical Facilities".

Partitioning and protecting network boundaries with firewalls

Besides the basic physical security of a site, the next most important aspect is controlling digital access into and out of the organization's network. In most cases this means controlling the points of connectivity to the outside world, typically the Internet. Almost every medium and large-scale company has a presence on the Internet and has an organizational network connected to it. In fact there is a large increase in the number of

©2005 American Power Conversion. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted, or stored in any retrieval system of any nature, without the written permission of the copyright owner. *www.apc.com* Rev 2005-0

smaller companies and homes getting full time Internet connectivity. Partitioning the boundary between the outside Internet and the internal intranet is a critical security piece. Sometimes the inside is referred to as the “trusted” side and the external Internet as the “un-trusted” side. As a generality this is all right, however, as will be described, this is not specific enough.

A firewall is a mechanism by which a controlled barrier is used to control network traffic into AND out of an organizational intranet. Firewalls are basically application specific routers. They run on dedicated embedded systems such as an internet appliance or they can be software programs running on a general server platform. In most cases these systems will have two network interfaces, one for the external network such as the Internet and one for the internal intranet side. The firewall process can tightly control what is allowed to traverse from one side to the other. Firewalls can range from being fairly simple to very complex. As with most aspects of security, deciding what type of firewall to use will depend upon factors such as traffic levels, services needing protection and the complexity of rules required. The greater the number of services that must be able to traverse the firewall the more complex the requirement becomes. The difficulty for firewalls is distinguishing between legitimate and illegitimate traffic.

What do firewalls protect against and what protection do they not provide? Firewalls are like a lot of things; if configured correctly they can be a reasonable form of protection from external threats including some denial of service (DOS) attacks. If not configured correctly they can be major security holes in an organization. The most basic protection a firewall provides is the ability to block network traffic to certain destinations. This includes both IP addresses and particular network service ports. A site that wishes to provide external access to a web server can restrict all traffic to port 80 (the standard http port). Usually this restriction will only be applied for traffic originating from the un-trusted side. Traffic from the trusted side is not restricted. All other traffic such as mail traffic, ftp, snmp, etc. would not be allowed across the firewall and into the intranet. An example of a simple firewall is shown in **Figure 2**.

Figure 2 *– Simple firewall to network*

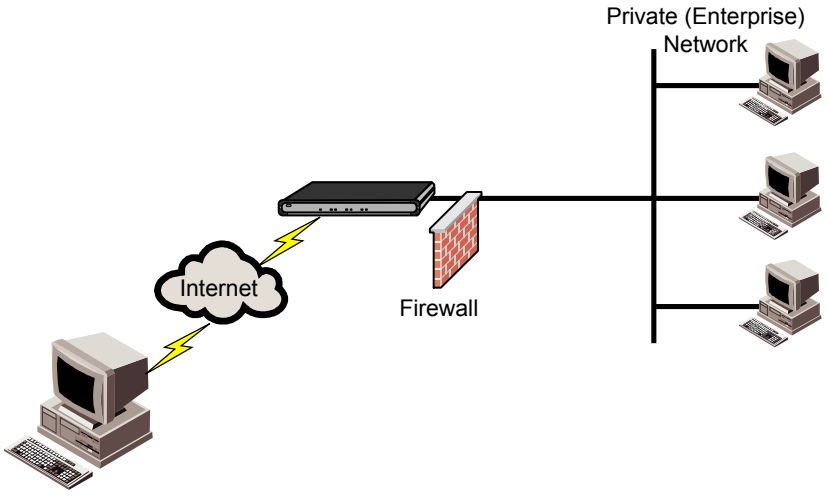

©2005 American Power Conversion. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted, or stored in any retrieval system of any nature, without the written permission of the copyright owner. *www.apc.com* Rev 2005-0

An even simpler case is a firewall often used by people with home or small business cable or DSL routers. Typically these firewalls are setup to restrict ALL external access and only allow services originating from the inside. A careful reader might realize that in neither of these cases is the firewall actually blocking all traffic from the outside. If that were the case how could one surf the web and retrieve web pages? What the firewall is doing is restricting connection requests from the outside. In the first case all connection requests from the inside are passed to the outside as well as all subsequent data transfer on that connection. From the exterior, only a connection request to the web server is allowed to complete and pass data, all others are blocked. The second case is more stringent as connections can only be made from the interior to the exterior.

More complex firewall rules can utilize what is called “stateful inspection” techniques. This approach adds to the basic port blocking approach by looking at traffic behaviors and sequences to detect spoof attacks and denial of service attacks. The more complex the rules, the greater the computing power of the firewall required.

One problem most organizations face is how to enable legitimate access to “public” services such as web, ftp and e-mail while maintaining tight security of the intranet. The typical approach is to form what is known as a DMZ (demilitarized zone), a euphemism from the cold war applied to the network. In this architecture there are two firewalls: one between the external network and the DMZ, and another between the DMZ and the internal network. All public servers are placed in the DMZ. With this setup, it is possible to have firewall rules which allow public access to the public servers but the interior firewall can restrict all incoming connections. By having the DMZ, the public servers are still provided more protection than if they were just placed outside a single firewall site. **Figure 3** illustrates the use of a DMZ.

Figure 3 *– Dual firewalls with DMZ*

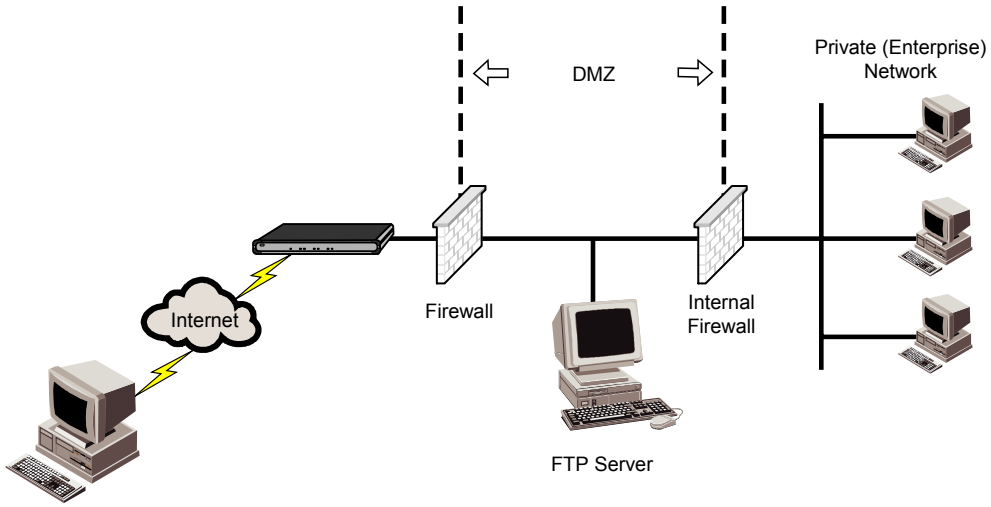

Using internal firewalls at various intranet boundaries can also help limit damage from internal threats and things like worms that have managed to traverse the border firewalls. These can even be run in standby so that normal traffic patterns are not blocked, but tight rules turned on in a problem situation.

©2005 American Power Conversion. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted, or stored in any retrieval system of any nature, without the written permission of the copyright owner. *www.apc.com* Rev 2005-0

Workstation firewalls

There is an important network security factor that most people are only now becoming aware of and that is that EVERY node or workstation on a network could be a potential security hole. In the past, basic attention was paid to firewalls and servers, however, with the advent of the web and the proliferation of new classes of nodes such as internet appliances, there are several more dimensions to protecting networks. A variety of worm virus programs hijack computers and use them to both further spread themselves as well as sometimes harm systems. Many of these worms would be stopped or greatly hindered if organizations had internal systems more "locked down". Workstation firewall products can block all port accesses into and out of individual hosts that are not part of the normal needs of the host. Additionally firewall rules on the INTERNAL side that block suspicious connections out of the organization can help prevent worms spreading back out of an organization. Between the two, both internal and external replication can be reduced. For the most part, all systems should be able to block all ports that are not required for use.

Basic Network Host Security

Port lockdown and minimizing running services

Many network devices and computer hosts startup network services by default, each of these services could represent an opportunity for attackers, worms and Trojans. Very often all of these default services are not needed. Doing port lockdown by turning off services reduces this exposure. As mentioned under the firewall section , similar to Network firewalls, desktops and servers can run basic firewall software to block access to unnecessary IP ports on the host or restrict access from certain hosts, This practice is important for internal protection when the outer defenses have been breached or from other internal threats There are many desktop firewall software packages available that do a great job of protecting hosts, for example, as of Windows XP Service Pack 2, Microsoft is actually bundling a basic firewall as well.

Username and password management

As mentioned in the introduction, poor username and password management is a typical problem in most enterprise networks. While sophisticated, centralized authentication systems (discussed later) can help reduce problems, there are basic guidelines that if followed can help tremendously. Four basic rules that need to be followed for usernames and passwords include:

1. Do not use obvious passwords such as spouse's name, favorite sports team, etc.
2. Use longer passwords with mixed numbers or symbols
3. Change passwords on a regular basis
4. NEVER leave default credentials in network equipment

Unless computers or equipment have built in policies that can enforce the above concepts, these are rules that must be self enforced. Rule (4) can at least be tested for by having network probes that try to detect equipment with default credentials.

©2005 American Power Conversion. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted, or stored in any retrieval system of any nature, without the written permission of the copyright owner. *www.apc.com* Rev 2005-0

Access control lists

Many types of equipment or hosts can be configured with access lists. These lists define hostnames or IP addresses that are valid for accessing the device in question. It is typical, for instance, to restrict access to network equipment from the inside of an organization's network. This would then protect against any type of access that might breach an external firewall. These types of access lists serve as an important last defense and can be quite powerful on some devices with different rules for different access protocols.

Securing Access to Devices and Systems

Since data networks cannot always be assumed to be protected from the possibility of intrusion or data "sniffing", protocols have been created to increase the security of attached network devices. In general there are two separate issues to be concerned about, authentication and non-disclosure (encryption). There are a variety of schemes and protocols to address these two requirements in secure systems and communication. The basics of authentication are discussed first and then encryption.

User authentication for network devices

Authentication is necessary when one wants to control access to network elements, in particular network infrastructure devices. Authentication has two sub concerns, general access authentication and functional authorization. General access is the means to control whether or not a particular user has ANY type of access right to the element in question. Usually we consider these in the form of a "User account". Authorization is concerned with individual user "rights". What, for example, can a user do once authenticated? Can they configure the device or only see data. **Table 2** is a summary of the major authentication protocols, their features, and their relevant applications.

Table 2 *– Summary of major authentication protocols*

Protocol	Features	Protocol Uses
Username \ Password	Plaintext, memorized token	Telnet, HTTP
CHAP (Challenge Handshake Authentication Protocol)	Uses hashes of passwords and time variant data to avoid straight password transmission	MS-CHAP, PPP, APC Http, Radius
RADIUS	CHAP or straight passwords, authorization and accounting methods	Backend for Telnet, SSH, SSL, Front end for Microsoft IAS Server. Typical central authentication method for network devices
TACACS+	Authentication, Authorization, Accounting, full encryption support	Cisco protocol, central authentication, some RAS use (Remote Access Service)
Kerberos	Service authentication and authorization, full encryption	Kerberized applications like telnet, Microsoft domain authentication service integrated with Active Directory

©2005 American Power Conversion. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted, or stored in any retrieval system of any nature, without the written permission of the copyright owner. *www.apc.com* Rev 2005-0

Restricting access to devices is one of the most important aspects of securing a network. Since infrastructure devices are supporting both the network and computing equipment ipso facto, compromising these can potentially bring down an entire network and its resources. Paradoxically, many IT departments go through great pains to protect servers, institute firewalls and secure access mechanisms, but leave some basic devices with rudimentary security.

At a minimum, all devices should have username password authentication with non-trivial (10 character, mixed alpha, numbers and symbols). Users should be restricted in both numbers and type of authorization. Care should be taken when using remote access methods that are not secure, i.e. usernames and passwords passed in the clear over the network. Passwords should also be changed with some reasonable frequency, perhaps every three months and when employees leave, if group passwords are used.

Centralized authentication methods

Appropriate authentication methods are important at a minimum, however, centralized authentication methods are even better when either a) large numbers of users for devices are involved or, b) large numbers of devices are in the network. Traditionally centralized authentication was used to solve problems found in situation (a); the most common was remote network access. In remote access systems such as dial-up RAS, the administration of users on the RAS network units themselves was just not possible. Potentially any user of the network could attempt to use any of the existing RAS access points. Placing all user information in all RAS units and then keeping that information up-to-date would exceed the abilities of RAS units in any large enterprise of users and be an administrative nightmare.

Centralized authentication systems such as RADIUS and Kerberos solve this problem by using centralized user account information that the RAS units, or other types of equipment, can access securely. These centralized schemes allow information to be stored in one place instead of many places. Instead of having to manage users on many devices, one location of user management can be used. If user information needs to be changed, such as a new password, one simple task can accomplish this. If a user leaves, the deletion of the user account prevents access for all equipment using centralized authentication. A typical problem with non-centralized authentication in larger networks is remembering to delete accounts in all places. Centralized authentication systems such as RADIUS can usually be seamlessly integrated with other user account management schemes such as Microsoft's Active Directory or LDAP directories. While these two directory systems are not themselves authentication systems, they are used as centralized account storage mechanisms. Most RADIUS servers can communicate with RAS or other network devices in the normal RADIUS protocol and then securely access account information stored in the directories. This is exactly what Microsoft's IAS server does to bridge RADIUS and Active Directory. This approach means that not only is centralized authentication being provided for the users of RAS and devices, but also the account information is unified with the Microsoft domain accounts. **Figure 4** shows a Windows Domain controller operating as both an Active Directory server and a RADIUS server for network elements to authenticate into an Active Directory domain.

©2005 American Power Conversion. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted, or stored in any retrieval system of any nature, without the written permission of the copyright owner. *www.apc.com* Rev 2005-0

Figure 4 *– Windows Domain Controller*

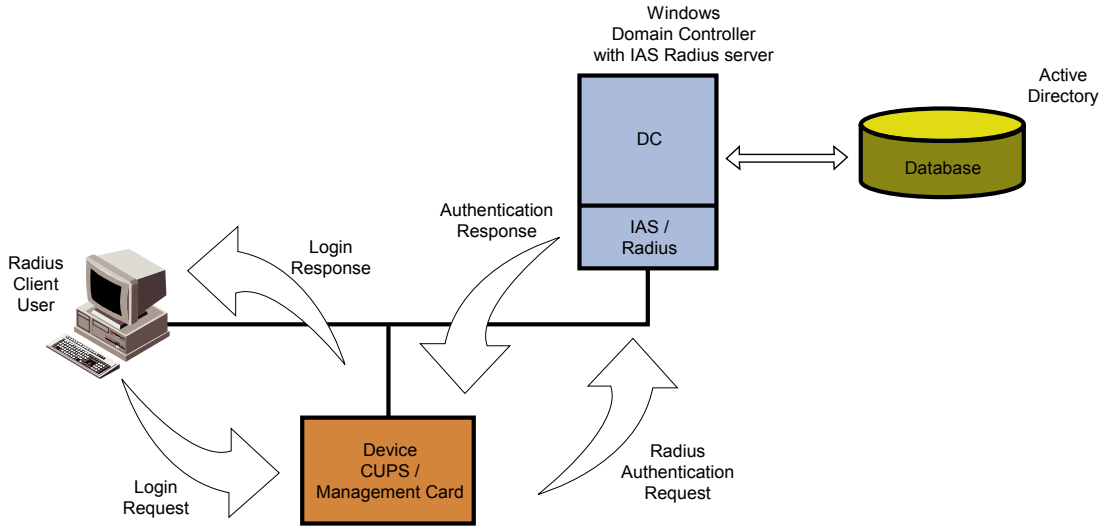

Securing network data with encryption and authentication

In some cases it is important to be concerned about disclosing information that is exchanged between network elements, computers or systems. Certainly it is not desirable that someone could access a bank account that is not theirs or capture personal information that may be transmitted over a network. When one wishes to avoid data disclosure over a network, encryption methods must be employed that make the transmitted data unreadable to someone who might somehow capture the data as it traverses a network. There are many methods to "encrypt" data and some of the major methods are described. With respect to network devices such as UPS systems, the concern is not traditionally about the value of protecting data such as UPS voltages and power strip currents; however, there is a concern with controlling access to these elements.

The non-disclosure of authentication credentials such as usernames and passwords is critical in any system where access is done over non-secure networks, the Internet for example. Even within organizations' private networks, protection of these credentials is a best practice. While it is less common, many organizations are starting to implement policies that ALL management traffic be secure (encrypted) not just authentication credentials. In either case, some form of cryptographic methods must be employed.

Encryption of data is usually accomplished by the combination of plaintext data (the input) with a secret key using a particular encryption algorithm (i.e. 3DES, AES, etc.). The result (output) is ciphertext. Unless someone (or a computer) has the secret key, they cannot convert the ciphertext back to plaintext. This basic methodology is at the core of any of the secure protocols (described later). Another basic building block of cryptographic systems is the "hash". Hash methods take some plaintext input and perhaps key input and then compute a large number called a hash. This number is a fixed length (number of bits) regardless of the size of the input. Unlike the encryption methods that are reversible, where one can go back to plaintext with the key, hashes are one way. It is not mathematically feasible to go from a hash back to plaintext. Hashes are used as special IDs in various protocol systems because they can provide a check mechanism on data

©2005 American Power Conversion. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted, or stored in any retrieval system of any nature, without the written permission of the copyright owner. *www.apc.com* Rev 2005-0

similar to a CRC (cyclic redundant check) on a disk file to detect data alteration. The hashes are used as a data authentication method (different than user authentication). Anyone trying to secretly alter data in transit across a network will alter the hash values thus causing detection. **Table 3** provides a basic comparison of cryptographic algorithms and their uses.

Table 3 *– Summary of major cryptographic algorithms*

Algorithm	Primary Use	Protocol Uses
DES	Encryption	SSH, SNMPv3, SSL/TLS
3DES	Encryption	SSH, SNMPv3, SSL/TLS
RC4	Encryption	SSL/TLS
Blowfish	Encryption	SSH
AES	Encryption	SSH, SSL/TLS
MD5	Hash, Message authentication codes	SSH, SNMPv3, SSL/TLS
SHA	Hash, Message authentication codes	SSH, SNMPv3, SSL/TLS

Secure Access Protocols

There are a variety of protocols such as SSH and SSL that employ various cryptographic mechanisms to provide security through authentication and encryption methods. The level of security provided is dependent upon many things such as the cryptographic methods used, the access to the transmitted data, algorithm key lengths, server and client implementations and most importantly, the human factor. The most ingenious crypto scheme is thwarted if a user's access credential, such as a password or certificate, is obtained by a third party. The classic case mentioned earlier is the password on a Post-It note on a person's monitor.

The SSH protocol

The Secure Shell (SSH) client-server protocol was developed in the mid 1990s in order to provide a secure mechanism to access computer consoles or shells remotely over unprotected or "non-secure" networks. The protocol provides "secure" methods by addressing user and server authentication and full encryption of all traffic exchanged between the client and server. The protocol has two versions, V1 and V2, which differ slightly in the cryptographic mechanisms provided. Additionally, V2 is superior in its ability to protect against certain types of "attacks". (An attempt by a "non-participating" third party to intercept, forge or otherwise alter exchanged data is considered an attack.)

While SSH has been used as a secure access protocol to computer consoles for years, it has traditionally been less employed in secondary infrastructure equipment such as UPS and HVAC equipment. However, since networks and the network infrastructure that support them are becoming more and more critical to the business practices of enterprises, using such a secure access method to all equipment is becoming more common.

©2005 American Power Conversion. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted, or stored in any retrieval system of any nature, without the written permission of the copyright owner. *www.apc.com* Rev 2005-0

The SSL\TLS protocol

While SSH has been the common secure protocol for console access for command-line like management, the Secure Socket Layer (SSL) and later the Transport Layer Security (TLS) protocol have become the standard method of securing web traffic and other protocols such as SMTP (mail). TLS is the most recent version of SSL and SSL is still commonly used interchangeably with the term TLS. SSL and SSH differ mostly with respect to the client and server authentication mechanisms built into the protocols. TLS was also accepted as an IETF (Internet Engineering Task Force) standard while SSH never became a full IETF standard even though it is very widely deployed as a draft standard. SSL is the secure protocol that protects http web traffic, also referred to as https for "http secure". Both Netscape and Internet Explorer support both SSL and TLS. When these protocols are used, a formal authentication of the server is made to the client in the form of a server certificate. Certificates are described subsequently. The client can also be authenticated with certificates, though usernames and passwords are most typically used. Because the SSL "sessions" are all encrypted, the authentication information and any data on web pages is secure. SSL is always used on web sites that wish to be secure for banking and other commercial purposes since clients usually access these sites over the public Internet.

Since web based management of network devices (embedded web servers) has become the most common method of basic configuration and point user access, protecting this management method is very important. Enterprises that wish to have all network management done securely, but still take advantage of graphical interfaces such as http, should use SSL based systems. As mentioned before, SSL can also protect other non-http communication. Should non-http based device clients be used, these systems should also employ SSL for their access protocols to insure security. Using SSL in all of these cases also has the advantage of using standard protocols with common authentication and encryption schemes.

Best Practices for Network Security

Well thought out security policies can significantly increase the security of a network. While policies can be both complex and cumbersome or basic and straight forward, it is often the simple aspects that prove most useful. Consider the combination of a centrally managed anti-virus update system and a host scanner to detect new or out of date systems. While this system would entail setup, central administration and software deployment capabilities, these are all generally available with today's Operating Systems. In general, policies and ideally automatic enforcement tools help reduce the obvious holes in system security so that one can concentrate on the more complex issues. The following would typically be part of an enterprise network security policy:

- Firewalls at all public-private network transit points
- Version controlled and centrally deployed firewall rule sets
- External resources placed in dual firewall, DMZ protected networks
- All network hosts lock down unneeded network ports, turn off unneeded services
- All network hosts include centrally managed anti-virus software

©2005 American Power Conversion. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted, or stored in any retrieval system of any nature, without the written permission of the copyright owner. *www.apc.com* Rev 2005-0

- All network hosts utilize central security updates
- Secure central authentication such as Radius, Windows/Kerberos/Active Directory
- Centrally managed user management with password policy (i.e. must change every three months and must be “secure password”
- Proactive network scanning for new hosts, out of date systems
- Network monitoring for suspicious behavior
- Incident response mechanisms (policies, manual, automated, etc.)

The above list represents the key items one should have in a policy. There are potentially other wide reaching items one could have in a policy. Of course, it’s always important to balance factors such as company size, risk analysis, cost and business impact when determining the type and breadth of a policy. As mentioned above, a system analysis is typically a good starting point, followed by the business analysis. While not obvious even very small companies should have some form of security policy, since all networks can be targets regardless of their size.

Conclusions

With the increased number of threats to networks such as worms, viruses and clever hackers, security can no longer be viewed as an option, even within “private” networks. Securing all equipment, including physical infrastructure equipment such as UPS systems and HVAC systems, is critical to maintaining uptime and seamless access to services. Providing and maintaining security across the enterprise typically means increased administration. Historically, this has been the largest barrier to broad implementations of security. Today, the amount of time spent repairing a network due to just a single worm or virus attack can easily be greater than the upfront time to more adequately secure an enterprise. Fortunately, there are many options in systems and software to increase the security of the network while reducing the overhead of managing such systems. Even basic practices such as periodic software updates, locking down all devices and using centralized authentication and secure access methods can go a long way to reducing risks. Institution of appropriate security policies and frequent network audits further increase the overall protection of the network.

About the Author:

Christopher Leidigh is the Director of Communications and Technology Research with APC based in Rhode Island, USA. He has 18 years of programming and design experience in computer and microprocessor systems with a current focus on communication systems, IP networking, security and embedded systems research and design. He has been with APC for ten years leading the embedded networking and management product line. He received a Bachelor of Science degree in Bio-Electrical Engineering from Brown University. He is a regular speaker at the Embedded Systems Conferences in the US and abroad as well as a member of the advisory board. Publications include the Journal of Physiology, Communications Systems Design, Embedded Systems Programming and the EE Times.

©2005 American Power Conversion. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted, or stored in any retrieval system of any nature, without the written permission of the copyright owner. *www.apc.com* Rev 2005-0