

Department of Homeland Security Cyberspace Security

Task Force 4

Network Security Equipment Deployment & Architectural Guidelines Working Group (5)

CONTRIBUTING AUTHORS

Victor Rychlicki,
*Working Group Chairman, Senior Security Systems Engineer,
Marconi Wireless*

Steve Macke,
Senior Consultant, Georgia Tech Research Institute

Marty Schulman,
Chief Technologist, Juniper Systems

Edward Roback,
Chief of Computer Security Division, NIST

Lawrence G. Dobranski,
Senior Network Security Architect, Nortel

Ron Mathis,
Director Security Infrastructure and Planning, Intrado

Greg Jackson,
VP, DeepNines Technologies

Table of Contents:

Executive Summary	4
Recommendation 1.0 - Securing Network Architectures	5
Recommendation 1.1	5
Recommendation 1.2	6
Recommendation 1.2.1	6
Recommendation 1.3	6
Recommendation 1.4	7
Recommendation 2.0 – Securing Cyberspace	8
Recommendation 2.1	8
Recommendation 2.2	8
Recommendation 2.2.1	9
Reference - Recommendation 1.x	
Introduction	10
Security is a Philosophy	11
The Human Firewall	11
Network Security-Specific Resources	12
Firewalls	12
Network Address Translation (NAT)	14
Authentication Systems	14
Intrusion Detection Systems (IDS)	15
Network Based	15
Host Based	16
Intrusion Prevention Systems (IPS)	16
Encryption	16
Host based / Gateway Anti-Virus	17
Network Security Auditing	17
Planning Security Domains and De-Militarized Zones (DMZ's)	18
Architectural Design	18
The Public DMZ	20
The Private DMZ	22
The Isolated Management Network	22
The Internal Network	23
Additional Architectural Security Options	23
Spoke and Hub 'Heavy' Security Architecture	24
Layered 'Light to Medium' Security Architecture	25
Layered 'Heavy' Security Architecture	26

Reference - Recommendation 2.x	
Introduction - Securing Cyberspace	27
Grading the Security of Individual Network Infrastructures	28
Security Classifications or 'DefCon' Levels	28
Level 1 Security Classification recommendations	29
Level 2 Security Classification Recommendations	29
Level 3 Security Classification Recommendations	30
Level 4 Security Classification Recommendations	30
Level 5 Security Classification Recommendations	31
Closing Statement	32
Additional Resources	34

Executive Summary

The focus of the five task forces created in December 2003 during the Cyber Security Summit in Santa Clara, CA. is to recommend and develop standards that will ultimately be implemented to heighten, or enhance the security of 'Cyberspace'. This is obviously a very large, complicated, and much needed endeavor.

As data infrastructures become more prevalent and widespread, much like our ever-increasing dependence upon them, it is necessary to ensure the security of these infrastructures so that they continue to deliver reliable and private transactions to all that are authorized to use them, and reject unauthorized transactions from those who are not authorized to use them.

The largest of these data infrastructures is of course, the Internet, otherwise referred to as "Cyberspace". Cyberspace continues to expand, and engulf more of our lives everyday with more and more of us relying on it to process financial and business transactions, as well as day-to-day personal email.

In order to fulfill the directive of the five task forces, securing cyberspace, a lot of very complex things must happen, and they must happen correctly since hackers, and would be attackers are very diligent, and detail oriented. It is the charter of this working group to point out two areas where it is believed that there is detail that needs to be addressed in the mission of securing cyberspace. These two areas are:

1. Recognizing that Cyberspace is made up of a collection of TCP/IP network infrastructures; and therefore planning security solutions using recommended equipment for each of these individual networks that ensures security; not at the 'box', or device level, but at the network level. This ensures that each individual network (that makes up a part of cyberspace), properly contributes to the overall security of cyberspace.
2. Recognizing the need for a method, or process for determining if the charter of the five task forces has been achieved. This means a method of determining, at any given time now or in the future, if cyberspace is actually secure; and at what level, or "Defense Condition" is this security actually providing.

It is the mission of this working group to recognize the need for these two areas, and also to provide starting points, and proposed recommendations in these two critical areas.

Recommendation 1.0 – Securing Network Architectures

It is recommended that the industry work together to develop a set of defined standards that make specific recommendations on using recommended security equipment, and best practices to understand, design, and implement secured IP network infrastructures.

This recommendation attempts to supply a starting point for the makings of new standards, and recommendations that give proper guidance to individuals and/or organizations on implementing recommended equipment into an IP network infrastructure in a secure manner.

Details of the specifics currently recommended as a starting point for this recommendation can be found in the “Reference” section of this paper under “Securing Network Architectures”.

This recommendation recognizes that work is being done to improve the process for certifying equipment thru organizations such as NIST with Common Criteria Certifications (CCC), and therefore improving the ability for certain types of IP network security equipment to be selected against a fair comparison of others for a particular security task, or tasks.

This recommendation also recognizes that there is a lack of standards and understanding on the proper use and implementation of recommended and selected security equipment.

This recommendation attempts to draw on the expertise of other Task Force Working Groups such as Task Force 4, Working Group 1, Common Configurations; and Task Force 4, Working Group 4, Best Practices for Technical Standards. Other input and recommended standards are also invited from any source including other standards organizations such as NIST, ITU, 3GPP, 3GPP2, ETSI, IETF, and TIA.

Details of the specifics currently recommended as a starting point for this recommendation can be found in the “Reference” section of this paper under “Securing Network Architectures”.

Recommendation 1.1 – Security is a Philosophy

It is recommended that discussion of this discussed topic be further developed and expanded so that industry may gain a more complete understanding of security, as this topic is more complex than many realize. Creating a completely secure IP network calls for more than good equipment and a secure architecture, it calls for an understanding of its users and its business practices.

Details of the specifics currently recommended as a starting point for this recommendation can be found in the “Reference” section of this paper under “Security is a Philosophy”.

Recommendation 1.2 – Network Security-Specific Resources

It is recommended that the following list of IP network security-specific resources be utilized and expanded upon as a reference for determining the types of equipment and functions required then implementing and planning a secure IP network infrastructure.

It is also recommended that this list be expanded upon and continually updated, as technology and security devices and services are continually being improved, and introduced.

- Firewalls
- Network Address Translation (NAT)
- Authentication Systems
- Intrusion Detection Systems (IDS)
 - Network Based
 - Host Based
- Intrusion Prevention Systems (IPS)
- Encryption
- Host Based / Gateway Anti-Virus

Details of the specifics currently recommended as a starting point for this recommendation can be found in the “Reference” section of this paper under “Network Security Specific Resources”.

Recommendation 1.2.1 – Network Security-Specific Resource Requirements

The devices selected in recommendation 1.3 by any given individual, or organization for security use in an IP infrastructure should meet all the requirements and be selected in accordance with the findings of Task Force 4 equipment recommendations, including basic system hardening and configuration as outlined by Task Force 4 Working Group 1 (Common Configurations).

Recommendation 1.3 – Network Security Auditing

It is recommended that a method or process for determining the importance of individual network devices, elements, services, and/or equipment be established so that a proper level of security can be established for each resource in the network infrastructure.

Details of the specifics currently recommended as a starting point for this recommendation can be found in the “Reference” section of this paper under “Network Security Auditing”.

Additional input on prioritizing systems, functions, or services for determining security requirements is required; as there are many different methods and viewpoints on what is considered important or critical. These viewpoints may be dependant upon the infrastructure owners’ use of the network.

Recommendation 1.4 – Planning Security Domains and De-Militarized Zones (DMZ’s)

It is recommended that all network architectures be planned and implemented in a secure manner. This recommendation attempts to suggest several possible layouts, or topologies that utilize recommended equipment, and practices to secure infrastructures against illegal or unintentional penetration, or attack. The use of techniques such as DMZ’s and restricted directional communications combined with the proper use of recommended equipment, and best practices can be used to create a secured IP architecture.

Details of the specifics currently recommended as a starting point for this recommendation can be found in the “Reference” section of this paper under “Planning Security Domains and De-Militarized Zones”.

It is also recommended that ongoing research in to these types of security deployment scenarios and their benefits become a regular event so that standards may be properly updated as security technology evolves. The industry as well as organizations, such as NIST, ITU, 3GPP, 3GPP2, ETSI, and TIA for example; should update relevant standards and proposed standards such as this one, to constantly include new security advancements in technology, or technology based security layouts and tactics.

Recommendation 2.0 – Securing Cyberspace

It is recommended that the industry work together to develop a defined set of standards that make specific recommendations on a method of determining the security level or security status of Cyberspace.

This recommendation recognizes that it is the charter of the five (5) task forces to put together standards and practices that facilitate the securing of cyberspace. There is currently no work on a method or process of determining how well these proposed standards and recommendations are working as they relate to the current or ongoing security of cyberspace, and attempts to offer a proposed recommendation for a solution.

This recommendation attempts to draw on the expertise of other Task Force Working Groups such as Task Force 4, Working Group 1, Common Configurations; Task Force 4, Working Group 4, Best Practices for Technical Standards, and Task Force 2, Early Warning. Other input and recommended standards are also invited from any source including other standards organizations such as NIST, ITU, 3GPP, 3GPP2, ETSI, IETF, and TIA.

Details of the specifics currently recommended as a starting point for this recommendation can be found in the “Reference” section of this paper under “Securing Cyberspace”.

Recommendation 2.1 – Grading the Security of Individual Network Infrastructures

It is recommended that an independent organization be created to test and rate, or grade the security and practices of individual networks. This independent organization may be responsible for defining, managing, updating, and executing this national testing and rating system for determining the ongoing security level of cyberspace.

Details of the specifics currently recommended as a starting point for this recommendation can be found in the “Reference” section of this paper under “Grading the Security of Individual Network Infrastructures”.

It is also recommended that additional work be done on the layout, makeup, and duties of the independent organization.

Recommendation 2.2 – Security Classifications or ‘DefCon’ Levels

It is recommended that a series of grades, or classifications, be developed that can be used to determine the level, or amount of security (Defense condition) an individual network has, or maintains. These security classifications state the tested networks ability to repel, or defend against illegal, or unwanted intrusions and attacks. These classifications are ultimately used to grade the overall level of security of cyberspace, since cyberspace is a collection of individual networks.

Currently there are five (5) proposed classification levels. Details of the specifics currently recommended as a starting point for this recommendation can be found in the “Reference” section of this paper under “Security Classifications or ‘DefCon’ Levels”.

Details of the specifics currently recommended as a starting point for this recommendation can be found in the “Reference” section of this paper under “Security Classifications or ‘DefCon’ Levels”.

It is also recommended that further development of these testing procedures be developed by organizations, such as NIST, ITU, 3GPP, 3GPP2, ETSI, and TIA for example; so that the resulting guidelines will be able to serve as a more complete model to determine the actual level of security of Cyberspace.

Recommendation 2.2.1 - Security Classifications or 'DefCon' Levels Requirements

It is recommended that these classification level recommendations include the guidelines and recommendations of Task Force 4, Working Group 1, Common Configurations; and Task Force 2, Early Warning where appropriate.

Details of the specifics currently recommended as a starting point for this recommendation can be found in the "Reference" section of this paper under "Security Classifications or 'DefCon' Levels".

Reference - Recommendation 1.x

Introduction

One of the purposes of this paper is to propose recommendations and guidelines in architecting secure network infrastructures in which recommended security equipment and components are deployed.

The initial task discussed here is to suggest initial seed ideas to be used by the security community to develop guidelines in architecting secure network infrastructures in which recommended security equipment and components are deployed. As these are developed, fleshed out, and subject to peer review, the resulting guidelines will be able to serve as a generic base reference model for creating a secure network topology. Our initial suggestions, which follow, are intended to be generic to most types of network topologies without respect to their individual purposes, functionalities, or intended operations.

Note: There are additional security precautions that may be put into place that may be dependent upon a network's specific purpose; for example, banking, education, and ISP specific networks. These purpose specific precautions, as well as protocol level security issues, are outside the scope of this paper.

Reference material on design and implementation of network architectures in a secure method is typically a difficult task and not readily available to the public. There is an abundance of information available on specific components used to secure networks, such as firewalls and Intrusion Detection Systems (IDS). There is ample information available, and best practices for these common systems to secure existing networks. The dilemma is that there is little availability of information on how to architect a communications infrastructure from the ground up so that the given infrastructure is inherently more secure from the start than a network that is backward engineered with firewalls and IDS added to plug holes that are later found after the infrastructure has been compromised.

This paper endeavors to educate the reader on how to engineer an IP communications infrastructure that, by its very nature of the architecture and layout, is inherently more secure than most networks today that have security added as an afterthought. Total security is something that must be designed into the very foundation of the network architecture and not something that is thrown on after the fact. This is not to say that existing networks cannot be secure, this is simply stating that there are architectural issues to consider in creating secure networks, and that these issues are applicable to both existing and planned infrastructures.

Following is discussion, and some basic architectural lessons and recommendations (in an informal whitepaper format) on two of the most commonly overlooked and least publicly understood issues facing network security. They are the 'Human Firewall' and secure architectural design.

Security is a Philosophy

One of the most important things to remember about security, all security, regardless of whether were securing a network infrastructure or a house, is that security is not a 'box' that you purchase, it's a philosophy, a way of doing things, and to some extent it's also a way of life. When securing anything, specifically a network infrastructure, it's important to know how to implement security practices and equipment (boxes). If implemented improperly, even the best designed security boxes, or devices, can be rendered useless.

The Human Firewall

In keeping with the idea that security is a philosophy we also have to understand that like philosophy, security comes in layers. When properly layered, each of these layers reinforces the others; and compromise of one layer does not compromise the other layers. One aspect, or layer of security that is overlooked is what is sometimes referred to as the 'Human Firewall'. Much like the hardware version of the firewall that protects networks today, the human firewall is also intended to block improper information flow.

This improper information flow can come in many forms, and is typically the result of human behavior. The most common forms of information leaks that result in information being revealed that attackers can use to penetrate networks come from:

Support web sites: Knowledge-base type articles that are intended to allow customers to service themselves with network access problems. These sources of information can sometimes provide key information such as IP addresses, vendor equipment types, and configurations that can give hackers just the bit of information they need to take advantage of a known, or even un-known weakness in the sought after infrastructure.

Phone based tech support personnel: The same issues apply here as with support web sites with some possible additional security issues. Phone based support personnel are in the business of helping people resolve some technical issue, a lot of the time these support personnel do not actually know the people they are talking to, this creates a personnel authentication issue. There have been many instances in corporate history where a hacker has simply placed a phone call to a corporate help desk acting as though he or she is a new employee, or newly hired consultant and for some reason has not yet been given access to the network and must have access to get a newly assigned project completed. In the spirit of assisting, as help desk do, access is often granted to someone unknown by assigning him or her a username and password. This is the ultimate form of hacking the 'human firewall'.

Press Releases: Often times a service provider, whether it's a Telco, or a bank offering new on line services, will make it publicly known that they have just finished implementing a public, or publicly accessible infrastructure. Often times as part of a joint marketing venture between the purchaser, and the supplying vendor, a press release will be sent out stating the availability of this new infrastructure and that the equipment is supplied by vendor X and utilizes the vendors new whiz-bang XXXX model products capable of offering XXX services. This is information that is valuable to hackers, as this information details the equipment used, how it's used in a general sense, and the fact that it's new brings the possibility that there are unknown and yet unfixed security vulnerabilities just waiting to be exploited.

These are just a few examples of ways that seemingly innocent information and people can be taken advantage of in ways that are not typically considered when laying out security policies and practices within an organization. As we mentioned before, security comes in layers, and the human layer is often the most important and the most overlooked.

Human security specifically, is outside the scope of this papers expertise. The above description, and discussion is supplied solely for information reasons to present one of the most commonly overlooked and important security vulnerabilities. 'Human Firewall' security will not be covered in any detail in the remainder of this paper.

Network Security-Specific Resources

The following is a listing of and definition of the most commonly used, and new IP networking security-specific tools, or resources. These resources are defined here, and later discussed as to their proper placement and use in securing most common IP network infrastructures.

Firewalls

Network Address Translation (NAT)
Authentication Systems
Intrusion Detection Systems (IDS)
 Network Based
 Host Based
Intrusion Prevention Systems (IPS)
Encryption
Host Based / Gateway Anti-Virus

Firewalls

A firewall is a combination of hardware and software used to implement a security policy governing the network traffic between two or more networks. A network firewall serves as a primary line of defense against external threats to an organization's computer systems, networks, and critical information. Because the firewall is a primary line of defense, the administration of this system must be carefully scrutinized. Segregation of duties, logging, auditing, and change control processes must be in place and continuously reviewed.

There are three firewall types. In order of increasing security, they are:

Packet Filtering Routers/Firewalls –Restricts network traffic by looking at the sources and destinations of individual network packets. There is no consideration of packet content or authorized use. Basic packet filtering can be implemented on many network routers, most commonly known as Access Control Lists (ACL). An enhancement to basic packet filtering is stateful inspection. This technology keeps tracks of “conversations” outside is in response to traffic from the inside.

Proxy/Circuit Level Gateway Firewalls –Acts as an intermediary for user requests at the connection level by requiring each user to first connect to the firewall. The firewall then establishes a second connection to the user’s final destination.

Application Proxy Firewalls –Extends the concept of Proxy/Circuit Level Gateway firewalls to the application level. Each application proxy inspects the traffic it is relaying to ensure that it conforms to that particular application’s protocol. For example, an FTP application proxy examines the user’s requests to verify conformance to the FTP protocol specification.

There are also hybrids of these technologies available, for example, a firewall that uses stateful inspection in combination with application proxies.

Proxy/circuit level gateways and application proxy firewalls may also require individual users to authenticate themselves (e.g., with a password) to the firewall before they establish the connection to the final destination. In addition, some firewall vendors offer hardened or secure operating systems as their firewall platform. These features can, in some circumstances, provide an even greater level of security and better protect the firewall should a security breach occur.

Regardless of the type of firewall selected, the controls it should implement are basically the same:

Permitted Services - The services allowed to traverse the firewall should be restricted to the smallest set required to implement the particular application or function. This restriction should be applied separately for each of the networks that the firewall interconnects.

Restricted Communications Flow - The direction of communications should be restricted and controlled between the networks the firewall interconnects. As a result, a clearly defined and limited communication trust model can be documented and monitored. For example, while it may be necessary for internal systems to initiate connections with a server on a DMZ network, it

should never be necessary for a server residing on a DMZ network to initiate connections with internal systems. Therefore, the firewall should not permit such connections. As a result, if an attack occurs, the scope of the attack is limited to the networks and systems controlled within this trust model.

Access Control - The particular set of systems or users allowed to use each service should also be restricted. For example, access to a database server on a DMZ network should be restricted to a) the web servers that retrieve information from the database and b) the internal system(s) used by the database administrator(s).

Control Messages –To make it more difficult to scan the firewall and determine what protocols may pass through it, the firewall should not return any protocol control messages such as “host unreachable,” “port unavailable,” “time exceeded,” etc.

Network Address Translation (NAT) – described below.

Network Address Translation (NAT)

NAT allows internal network topology and addressing to be hidden from external users by using one set of addresses to access the external network, a different set of addresses to access the internal network, and a mapping between the two.

Authentication Systems

Authentication provides a means for identifying an object (e.g., user, application, system, etc.). As a result, the object can then be granted access to only those services it requires and its activities can be monitored. A variety of authentication mechanisms are available, ranging from simple password-based systems to token-based to biometrics. The particular authentication technology selected is dependent upon the classification assigned to the data, system, or network. For example, as a primary line of defense, a firewall would be classified as critical. Therefore, minimally, a token-based system should be used to authenticate with administrator privileges to the firewall.

In addition, an organization may extend the authentication mechanism to ensure that a particular transaction(s) can be traced back to a particular user (e.g., brokerage transactions). This is commonly known as providing non-repudiation capabilities.

The most common authentication protocols used are:

- Remote Authentication Dial-in User Service (RADIUS)
- Kerberos
- Lightweight Directory Access Protocol (LDAP)
- Windows NT Domain Services
- Novell EDS
- Windows Active Directory

Intrusion Detection Systems (IDS)

Intrusion detection systems search for signs of unauthorized access or use. Network-based intrusion detection examines the types and contents of network packets; host-based intrusion detection examines system audit trails and activity logs. Unauthorized access and use can be detected in one of two ways: misuse detection searches for known attack “signatures,” much in the same way that anti-virus software searches for viruses; anomaly detection searches for unusual behavior based on profiles of expected user and application activity. Network-based IDS has the advantage of being able to monitor and alert on attacks of all systems on an entire network segment. This capability generally makes network-based IDS easier and less expensive to deploy. However, network-based IDS’s are limited because it cannot “see” what is happening on individual hosts. For this reason, a complete implementation will make use of both network-based and host-based solutions.

Network-Based IDS

An intrusion detection system should be deployed on each DMZ network, extranet segments, as well as on the internal network segment that is connected to the firewall. Optionally, an intrusion detection system may also be deployed on the “Internet side” of the firewall; however, this system must be carefully configured to avoid unnecessary alarms.

In all cases, the IDS should be configured with two network interfaces: one for receiving traffic to be analyzed and the other for reporting alarms to an IDS management console(s). The interface used for receiving traffic should be configured without a network address (in what is known as “stealth mode”), making it almost impossible for attackers to identify its location or existence.

For intrusion detection systems that are monitoring the activity of external networks (e.g., DMZs), the reporting interface should be connected to an isolated IDS segment, and communication between the IDS and IDS management console(s) must be controlled by the firewall. This approach has several benefits and should hold true with all security systems management interfaces. First, it uses the firewall to restrict access to the intrusion detection systems. Secondly, it improves the performance of the IDS’ analysis and reporting functions by segregating these duties between the two interfaces. Thirdly, keeping IDS traffic on a separate network avoids any reduction in available bandwidth afforded to the Public and Private DMZs. Finally, because the IDS traffic is flowing on separate network, it will not be interrupted should a bandwidth-consuming denial of service attack be directed at the Public DMZ.

Like the IDS systems deployed for external network monitoring, intrusion detection systems that are monitoring internal networks require segregation of receiving and transmitting interfaces. An organization can use the same architecture as described above, if desired. An alternative is to have the reporting interface connected to the internal network, but internal switching and routing would restrict access.

In either case, the communication between the IDS and the management console(s) must be strongly authenticated and encrypted. In addition, the system clocks of all systems that are monitored or play a role in the monitoring of intrusions (e.g., IDS, management consoles, firewalls, routers, DMZ systems, etc.) should be synchronized to a common time to allow for correlation and auditability of log data from multiple systems. This configuration may require the installation of an additional network interface on the firewall and/or additional firewalls/ routers.

Host-Based IDS

Host based IDS should be deployed on all critical systems. In order for host-based intrusion detection to function effectively, these systems must be configured to enable full auditing and activity logging. This may require that the systems be configured with additional memory and/or disk space to avoid adversely impacting their performance. All IDS data and management must be strongly encrypted and authenticated.

Host based security specifically is outside the scope of this papers expertise. The above brief description is supplied solely for information reasons to present all usable resources. Host based security will not be covered in any detail in the remainder of this paper.

Intrusion Prevention Systems (IPS)

Intrusion Prevention is quickly becoming the newest security resource, and the 'next generation' firewall technology so to speak. Intrusion prevention, in essence, combines the abilities of firewalls and IDS. IPS systems typically sit in-line in the main traffic stream in a similar fashion to firewalls. IPS's perform most, if not all functions a firewall performs as far as managing traffic that flows through it. Intrusion prevention takes firewalling a step farther by also being capable of mitigating several types of intrusions or penetrations attempts that standard firewalls cannot detect. Also, IPS systems look for incorrect, or non-typical behavior in traffic patterns and take action appropriately to protect the network.

Some IPS systems are designed in a manner that allow them to be placed outside the network they are protecting in a manner that does not interfere with traffic routing, or addressing, since the IPS device is not detectable in the infrastructure and does not participate in network activities the way a normal firewall or proxy device would.

Encryption

Encryption is an essential part of security and should be implemented wherever confidentiality is a concern. However, in many cases, encryption only protects the data while it is in transit. A more secure implementation would also encrypt the data once it is stored at its final destination.

Encryption specifically is outside the scope of this papers expertise. The above brief description is supplied solely for information reasons to present all usable resources. Encryption based security will not be covered in any detail in the remainder of this paper.

Host based / Gateway Anti-Virus

Anti-Virus software should be implemented on all desktops or email servers and should be updated regularly for the latest signatures. Additionally, given the number of users within any network that have external POP email addresses, such as Yahoo or Hotmail; organizations have begun to deploy gateway anti-virus solutions that could otherwise bypass the email server.

Anti-Virus, specifically email security, is outside of the scope of this paper. The above brief description is supplied solely for information reasons to present all usable resources. Anti-Virus will not be covered in any detail in the remainder of this paper.

Network Security Auditing

The first step in creating a secure IP architecture is to know and understand what you're securing. This may sound trivial, but you can't secure what you don't understand. Every *resource* that is a part of, or is connected to, the network must be analyzed and categorized so that they may be properly laid-out, accounted for, and architected into the security scheme.

Prioritizing Systems

These network resources, as mentioned above, may follow under the headings of:

Servers

- Web based publicly accessible (i.e., extranet, services, etc.)

- Web based internally accessible (i.e., intranet)

Application based publicly accessible (i.e., on-line banking)
Application based internally accessible (i.e., Corporate email)

Databases

Publicly accessible (i.e., on-line bill retrieval)
Internally accessible (i.e., customer database)

Network Management Systems (NMS) (incl. servers/databases, non-security related)

Networking Equipment (switches, routers, etc., non-security related)

The list of resources generated and categorized by a close examination of each networks individual use, or service requirements should then be cross-referenced by three additional topics:

Accessibility – How accessible does this specific resource need to be, and who will be accessing it? Is this resource to be publicly accessible, or only privately accessible?

Functionality – What purpose does this resource serve? How critical is it to the overall service offering of the network infrastructure?

Vulnerability – How vulnerable is this resource? Is this vulnerability acceptable given its function? How replaceable is this resource?

Planning Security Domains and De-Militarized Zones (DMZ's)

The second step in creating a secure IP architecture is plan out your security zones. This is done using the matrix created by the network resource audit described above with modifications for geographical, departmental, financial, and/or political issues.

In most situations one DMZ, or security domain, should be implemented in the network infrastructure. This DMZ should be positioned between the trusted (internal) network and the untrusted (typically the Internet) network. Typically this initial DMZ will contain only resources that are required to be publicly available, such as web servers, mail servers, and domain name servers (DNS).

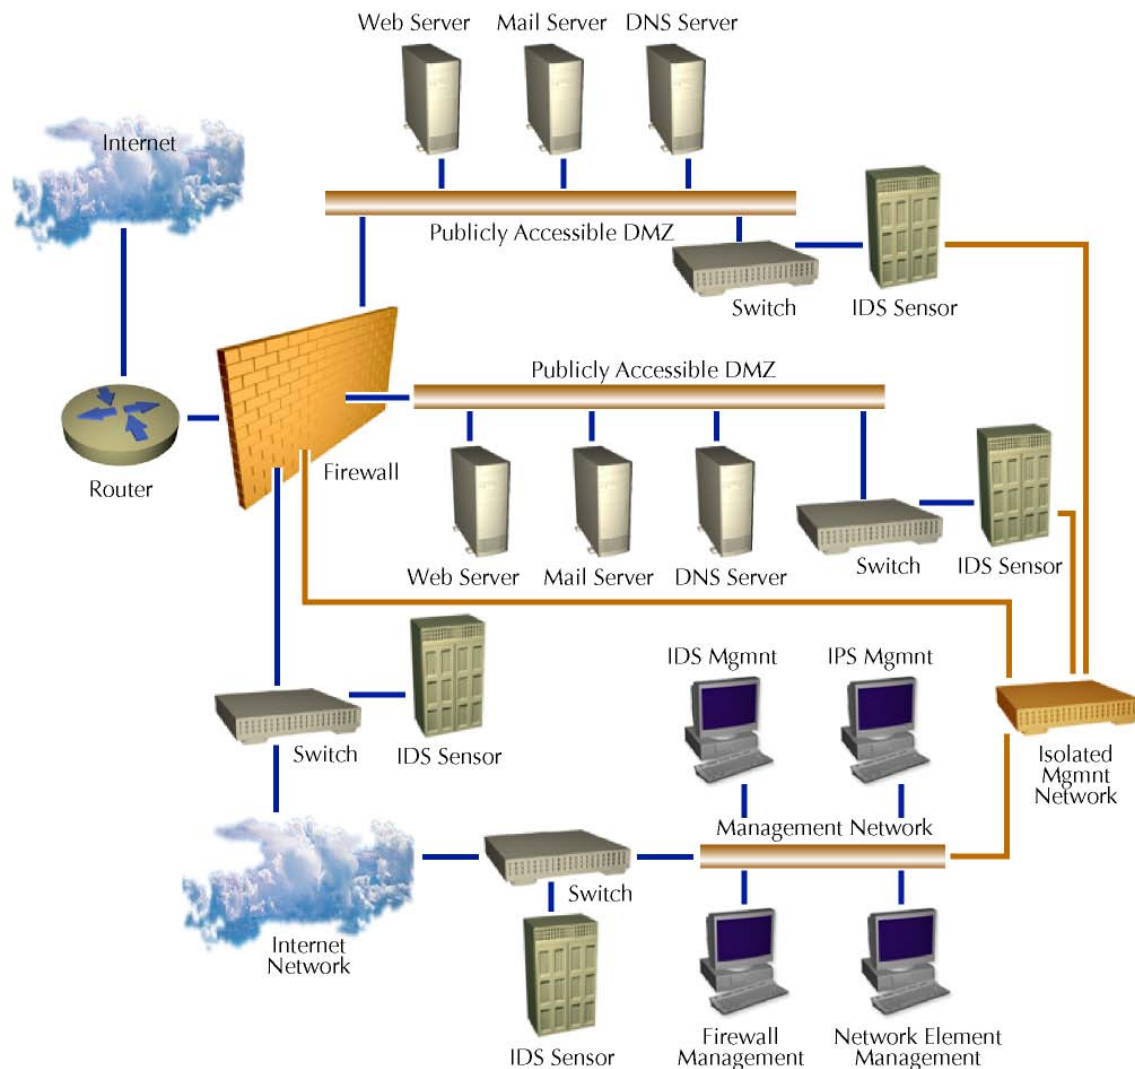
Given the specific functions and services that may be offered and required by different networks, it may be desirable to create two or even three security domains to provide maximum protection to all resources. Each set of systems or applications that require different levels of accessibility, or access requirements should be implemented in its own DMZ or security domain. Subnets work best to separate different security zones with security-specific equipment demarking each zone. Each security zone may also be set up with individual subnets according to function. Having FTP servers on one subnet and restricting that subnet to FTP traffic only, and web servers on another subnet while restricting that subnet to HTTP traffic only, in the same DMZ is a good way to be sure that any FTP weaknesses aren't used to exploit your web servers.

Another suggestion is to place similar equipment or resources on common switches divided by virtual local area networks (VLAN). For example, place untrusted subnets or zones on the untrusted switch, DMZ resources on the DMZ switch, and trusted internal resources on the internal switch. Cabling from each of these specific switches, or VLANS should be running to security specific demarcation equipment such as firewalls, or Intrusion Prevention Systems (IPS).

Architectural Design

The reference diagram below details one possible architectural deployment scenario designed to provide a light to moderate level of network security. Note that there are many other areas that can be addressed to enhance security such as deploying Host Based Intrusion Detection, Virus Scanning gateways, and proxy servers. These additional security measures, or layers, are

outside the scope of this paper as we are concentrating on the secure architectural layout of networking infrastructures at a higher level.



This reference network does not take into consideration equipment redundancy, or additional equipment for load sharing, or performance issues. Adjustments may be made accordingly.

Note: This and the following diagrams also show two mail servers. Two mail servers are not normally the case in a standard network deployment. The placement of the actual mail server in real deployment would be at the discretion of the security architect. Since the mail server serves as both a publicly and a privately accessible device, placement of the mail server in one DMZ or the other is not mandatory, but must reside within a DMZ. One thought is that by placing the mail server in the private DMZ, the internal company could continue to run if the public DMZ was compromised and brought down by an attack or other destructive event.

This network employs a standard firewall for external protection between the trusted and untrusted (Internet) network. IDS sensors are used to monitor and detect attacks and improper traffic behavior should an attack make it past the firewall or if an attack should be generated from inside the trusted network. Each of these security zones is attached to an individual port of the main firewall.

Three different security zones are configured in this scenario:

A publicly accessible DMZ
A Private DMZ
Internal network / Management Network

Connections from hosts from the Internet are controlled as follows:

Hosts on the Internet may initiate connections to the servers on the Public DMZ using only those protocols needed by the particular application(s) offered. For example, the Web server may only be reached with valid HTTP and SSL requests, the mail server may only be reached with valid SMTP requests, and the Name server may only be reached with valid DNS queries.

Access to any other network in the architecture is prohibited.

Hosts on the Public DMZ are permitted to initiate connections to hosts on the Private DMZ. Again, these connections are restricted to only the particular hosts and protocols required.

Hosts on the internal network may initiate connections to the servers on the Public DMZ using required protocols and the protocols necessary to administer the servers. Access using administrative protocols is limited to the particular systems used to perform administrative functions (e.g., the Web server may only be accessed by the Web administrator's system, the mail server may only be accessed by the mail administrator's system, etc.)

No connections may be initiated from this network directed to the internal network.

Connections to and from the Private DMZ network are controlled as follows:

No connections may be initiated from the Internet/extranet to the Private DMZ.

Hosts on the Public DMZ are permitted to initiate connections to hosts on the Private DMZ. These connections are restricted to only the particular hosts and required protocols.

Hosts on the internal network may only initiate connections to the private DMZ for administrative purposes. Access using administrative protocols is limited to the particular systems used to perform administrative functions (e.g., the database administrator is permitted to perform maintenance and administration and the primary database server can perform data uploads and downloads).

No connections may be initiated from this network directed to the internal network.

Connections to and from the isolated Management network are controlled as follows:

No connections may be initiated from the Internet, Public DMZ, or the Private DMZ to the isolated IDS segment. However, if the monitoring of the IDS systems is outsourced, authenticated, encrypted, and restricted to a specified IDS management console, then connections may be allowed and initiated from the outsource company to the isolated segment.

Depending on the alerting mechanism, the isolated IDS segment may initiate connections with the Public DMZ (e.g., the mail server) for the purpose of intrusion alerts (e.g., email/ pager).

No connections may be initiated from this network directed to the internal network.

Connections to and from the internal network are controlled as follows:

No external network (including the Internet, Public DMZ, Private DMZ, or isolated IDS segment) may initiate connections directed to the internal network.

The internal network may initiate connections to any external network. However, this access should be limited to the particular systems and protocols necessary to perform a specified function.

The firewall should perform Network Address Translation (NAT) for all networks it inter-connects. Other than as described in this section, all connections traversing the firewall are prohibited using the firewall strategy commonly referred to as “that which is not expressly permitted, is denied.”

The Public DMZ

Each application provided to the public is implemented on a separate server attached to the DMZ. This segregation of duties limits the amount of damage and disruption that can be caused by a security breach. It also allows the configuration of the privileges on each server to be as restrictive as possible. This includes configuring each server to only accept connections that are required for the particular application they offer.

Each server on the Public DMZ should implement full auditing and activity logging. If possible, a dedicated log server on the internal network should retrieve logs from these systems in a timely manner. This log server protects the logs from unauthorized modification and/or deletion. In addition, each host should run system-level vulnerability assessment and host-based intrusion detection applications. The data generated by these applications should be encrypted to ensure confidentiality and employ strong authentication to mitigate the risk of forgery.

Hosts on the Public DMZ should not be permitted to establish connections to the internal network. Therefore, all system administration, security management activities, and data transfers should be initiated from the internal network. This includes, but is not limited to, Web content changes, invocation of assessment tools, retrieving assessment data, retrieving log data, and host-based intrusion alarms.

Because the servers on the Public DMZ are accessible to the public, it should be assumed that they are likely to be attacked. Therefore, these systems should only contain data that originates from systems not on the Public DMZ.

The Private DMZ

Each application provided to the hosts on the Public DMZ must be implemented on a separate server attached to the Private DMZ. This segregation of duties limits the amount of damage and disruption that can be caused by a security breach. It also allows the configuration of the privileges on each server to be as restrictive as possible. This segregation includes configuring each server to only accept connections that are required for the particular application they offer.

Each server on the Private DMZ should implement full auditing and activity logging. If possible, a dedicated log server on the internal network should retrieve logs from these systems in a timely manner. This log server protects the logs from unauthorized modification and/or deletion. In addition, each host should run system-level vulnerability assessment and host-based intrusion detection applications. The data generated by these applications should be encrypted to ensure confidentiality and employ strong authentication to mitigate the risk of forgery.

Hosts on the Private DMZ should not be permitted to establish connections to the internal network. Therefore, all system administration, security management activities and data transfers should be initiated from the internal network. This includes, but is not limited to, database changes, invocation of assessment tools, retrieving assessment data, retrieving log data, and host-based intrusion alarms.

The Isolated Management Network

The isolated segment provides for secure and controlled communication between all security systems and its associated management console. Essentially, this is an additional external network that is connected and controlled independently by security administrators, accessed only by security administrators.

Each IDS or security device should be configured with at least two network interfaces: one for receiving traffic to be analyzed and the other for reporting alarms to a management console(s). The interface used for receiving traffic should be configured without a network address ("stealth mode"), making it almost impossible for attackers to identify its location or existence.

The reporting interface should be connected to this isolated IDS segment for the purpose of communicating only with the associated management console. This communication should employ strong authentication and encryption (e.g., point-to-point encryption with public-private key authentication) to mitigate the risk of forgery. This protection has several benefits. First, it allows the act of segmenting the network to restrict access to the IDS, allowing communication from only the associated management console(s). Secondly, it improves the performance of the IDS' analysis and reporting functions by segregating these duties between the two interfaces. Thirdly, keeping IDS traffic on a separate network avoids consuming bandwidth in the Public and Private DMZs. Finally, because the IDS traffic is flowing on separate network, it will not be interrupted should a bandwidth-consuming denial of service attack be directed at the Public DMZ.

The Internal Network

The firewall should restrict all external systems (e.g., the DMZ networks and the isolated IDS segment) from initiating connections with the internal network. The internal network may initiate connections with these external networks as described in previous sections. Additionally, a network-based intrusion detection sensor should also be deployed on the internal network segment that is connected to the firewall (also with a stealth mode interface). This sensor provides "last-resort" monitoring for any inappropriate traffic entering through the firewall, as well as for traffic leaving the internal network. A filtering router should be used to limit access to the sensor's alarms/administration interface; only the IDS Management console(s) should have access to the sensor.

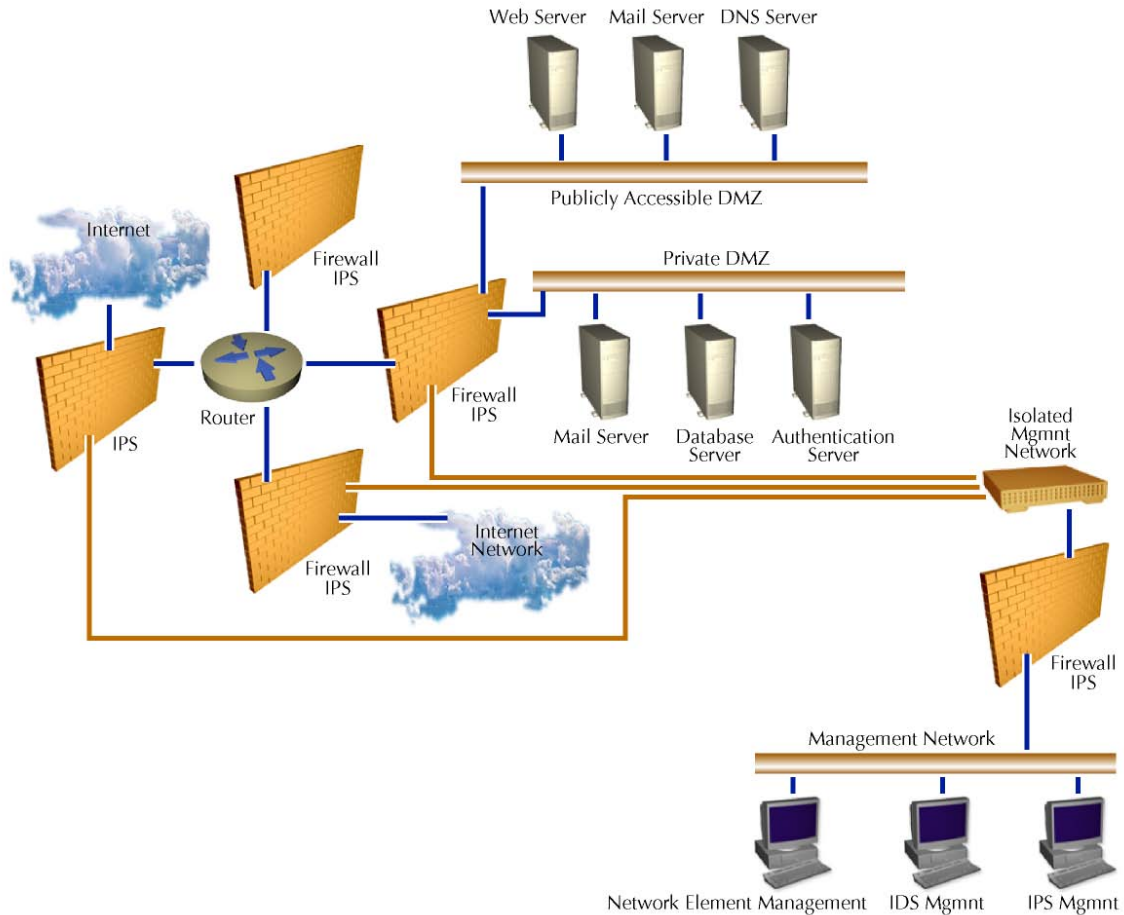
Within the internal network, all IDS and Security Management Consoles should reside on a dedicated network segment(s). Access to this segment should be restricted to appropriate personnel/systems using firewalls or filtering routers and strong authentication.

Additional Architectural Security Options

Below are some additional architectural diagrams displaying other options for architecting networks in a secure manner in the order of increasing security. The above diagram is representative of a star, or spoke and hub configuration, of laying out a secure network infrastructure. Other options such as the 'layered' approach are also available.

Inherently the 'layered' approach tends to be more secure, as the infrastructure is created in layers that must be penetrated in order to reach the internal network segment, as opposed to the spoke and hub approach which places each network segment at equal distance from the untrusted perimeter.

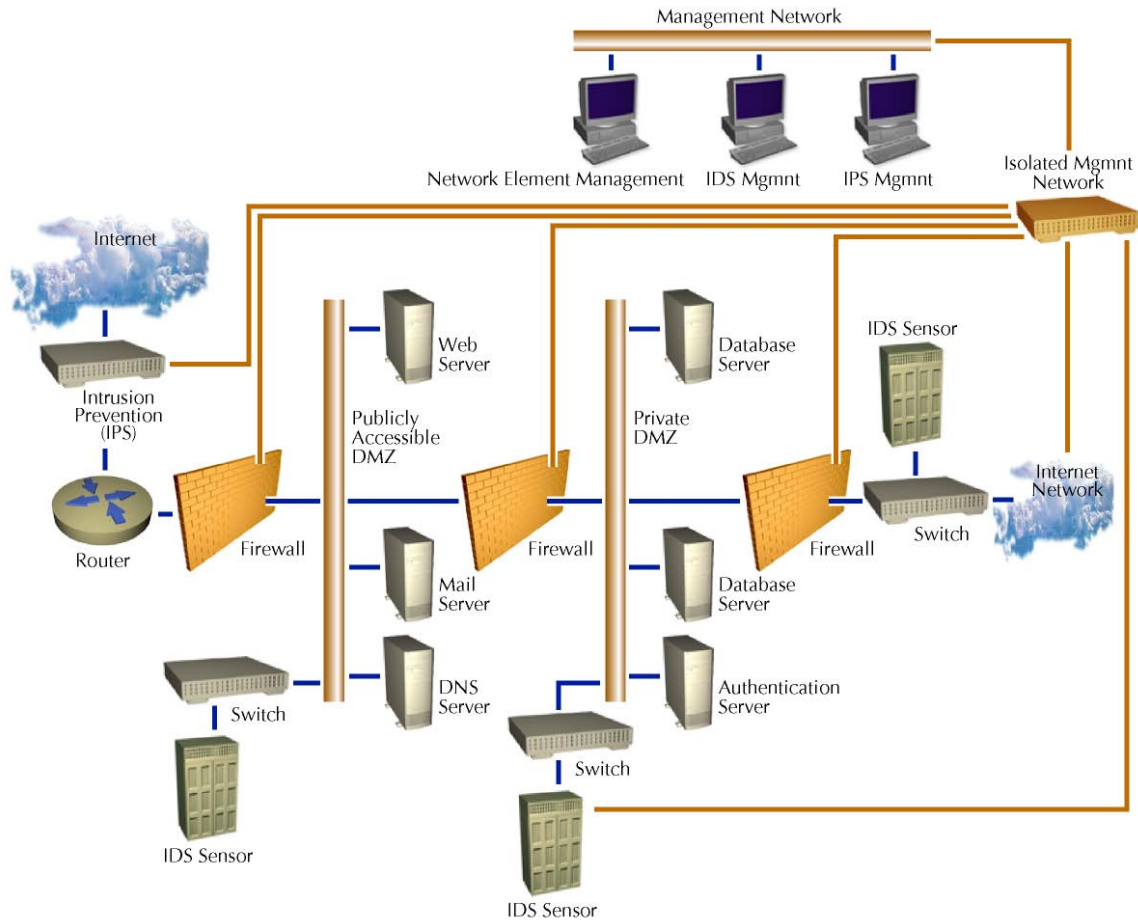
Spoke and Hub 'Heavy' Security Architecture



This star, or spoke and hub approach, is built in a manner that offers heavy security as opposed to the previously shown configuration that offers light to medium security. The major differences here are that Intrusion Prevention Systems (IPS) may be utilized with, or in place of Firewalls. Some IPS's may also be deployed 'outside' the network in front of the main router and they operate in-line invisible, and undetectable.

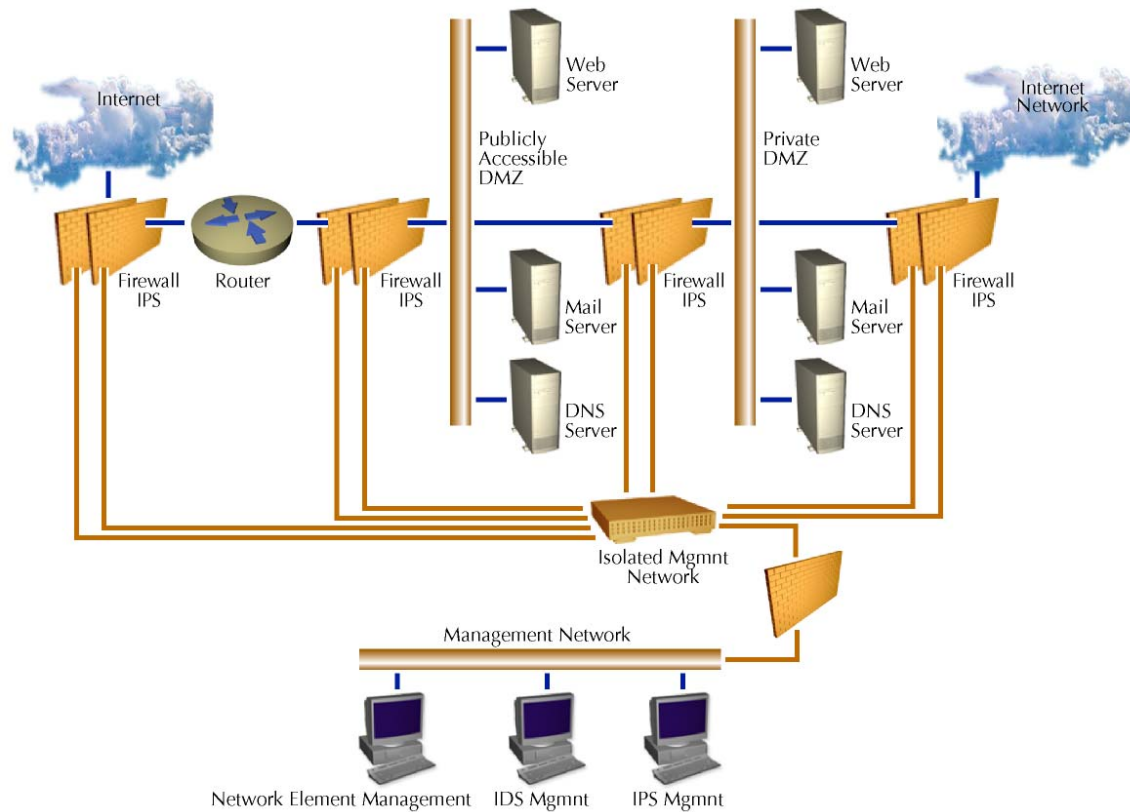
Inside the network, additional IPS's or firewalls may be used to segment off each security zone independently. Also, note that there is no direct connection from the internal network to the management network, and also note the additional firewall or IPS that is present on the management network. This created additional security for the all-important management network as it is now completely stand-alone, and contains it's own protection.

Layered 'Light to Medium' Security Architecture



The above 'layered' security approach is inherently more secure than the spoke and hub architecture in that it places each network in-line and therefore the most sensitive network is protected behind one or more additional networks. Instead of each network being placed equal distance from the untrusted perimeter as with spoke and hub architectures, the more sensitive, or important networks, are hidden behind the other less important networks.

Layered 'Heavy' Security Architecture



This layered approach may be considered the most secure of all shown approaches. Not only does this approach utilize the layering method of security architecture, but also it utilizes multiple firewall or IDS systems at each segment or security boundary.

The best approach to this method is to use multiple devices, one firewall and one IPS, or two IPS's at each boundary with each device at the given boundary utilizing a different operating system, and security software load. This method ensures that a security weaknesses in one operating system, or security software load, will not likely not be an issue in the next device in line since the next device will utilize a different operating system, and security software load.

As a final note, be sure that all unused and unneeded ports and services in each security device is disabled, and/or removed or uninstalled. Also, security systems, like all other systems, should always be kept up to date with the latest software version and patches. These actions keep the devices themselves more protected and 'hardened' against weakness and attacks.

Reference - Recommendation 2.x

Introduction - Securing Cyberspace

The goal of this, and the cumulative work of the other DHS Task Forces and working groups is to secure cyberspace, in relation to the boundaries of the United States. This is indeed a monumental task. The aim of this working group is to educate and recommend network level architectural issues that can lead to this noble purpose

Essentially, the Internet, or 'Cyberspace', is NOT a collection of boxes performing individual tasks. Cyberspace IS a collection of networks performing functions, and should be treated as such. Cyberspace is essentially one large network made up of smaller individual networks, each adding to the whole. Each network, no matter how large or small, plays a part in making up Cyberspace. Since each of these networks are individually owned and operated, they cannot be secured as a whole against outside attack, especially since some of the attacks are likely to originate from somewhere inside the network (as a whole). Therefore, we are faced with protecting these individual networks that make up Cyberspace, from other networks inside Cyberspace, as well as protecting Cyberspace (as it exist in the USA) from malicious activity outside North America. So, we must create a plan for protecting Cyberspace one network at a time.

Any engineer will tell you that each network is different, and that is a very accurate statement. There are however, commonalities to all networks regardless of their purpose, be it on-line banking, Internet Service providers, or a simple home network. As far as security is concerned, there are also commonalities; the most common security device is the traditional firewall. Firewalls can be, and are implemented in just about every type of network ever built, even the small home network. There are also other issues that are typically overlooked that are common, or should be common, to all networks, one of these issues is secure design, meaning a network architecture that was built, or re-built with security, as well as function in mind.

Most networks today were not originally designed with security as one of the top issues; network security for the most part had traditionally been an afterthought. There is good reason for this. It's only been in the last decade or two that serious security threats have begun to commonly arise in cyberspace and many network owners have scrambled to install firewalls, Intrusion Detection, and other associated devices to plug these previously unknown security weaknesses.

The industry has recently discovered the need for enhanced security, and has invested billions in hardening its architectures. One thing that is often overlooked in this process, and in the process of designing new networks is that the layout of the architecture itself needs to be secure. In other words, a well laid-out network with no security devices can be inherently more secure than a poorly laid out network with an abundance of security devices. Therefore a few minutes spent learning how to securely architect networks would be time well spent, and this is what part of this paper aims to do.

If we, as an industry, or as a country, are to claim that our portion of Cyberspace is secure, then we must have some method of making this determination. We need some grading format that is common to all networks, that allows us to grade the security of the individual networks so that we can determine the overall security of the U.S. portion of Cyberspace, as our President has required us to do.

Grading the Security of Individual Network Infrastructures

One suggestion or recommendation for grading the security level of each individual network infrastructure is to implement a national and industry accepted rating system that is generic enough to encompass all networks, but detailed enough to ensure their security on many different levels. After all, security is a philosophy, and a way of life, or a way of doing business in this case; and any grading or rating system should take this into account.

The following proposed recommended rating system divides ratings, or certifications, up into five levels. Level 1 thru level 5 certifications, based on tested infrastructure security. These certifications can be analogized to the U.S. Military's DefCon 1-5 rating system for defense condition. For example, a network with a security classification of 5 has been tested and deemed more secure than a network with a security classification of 1. These certifications state not only the graded networks 'Defense Condition', but also the defense condition of the operator, and their business and response practices as it relates to their networks security. This certification also demonstrates the operator's ability to accurately and timely communicate information to the rest of the industry on lessons learned from dealing with IP or network infrastructure security issues.

In order to properly implement a system like this, it is recommended that an independent organization be created that defines, manages, updates, and executes, the delivery of these recommendations. This independent organization should be made up of industry and government security experts at all levels from Firewall/device experts, management level subject matter experts on incident response, and technical support directors, just to name a few. This independent organization should also have the ability and authority to grant certifications to other organizations thereby enabling them to conduct testing and grading on behalf of the independent organization.

Security Classifications or 'DefCon' Levels

The following five classification levels provide a means for rating, or grading an individual infrastructures security level. The word 'infrastructure' in this context, refers to the network, and the network owners security practices as it relates to the network and the information it contains.

The security classifications are listed and discussed in the order of increasing security requirements.

Level 1 Security Classification recommendations:

The level 1 classification of network security is the first rating and has the most basic requirements of the five. They are:

The network to be certified undergoes an external vulnerability assessment of all perimeter systems. These systems may include Internet gateways, web servers, email servers, name servers, modem dial-in connectivity, etc. This means that all systems or network components that are exposed to any other network not owned by the same owner must be tested.

The results of the testing must not return any vulnerabilities that are exploitable.

The vulnerability assessment must conclude that it is not possible to 'see', or identify components or systems, inside the tested network any farther than its security perimeter.

Vulnerability assessment to be carried out using the most up-to-date recognized industry assessment tools that are certified by the independent organization.

Other details of these level 1 recommendations are to be addressed by the independent organization.

Level 2 Security Classification Recommendations:

The level 2 security classification includes all the requirements of the level 1 classification plus:

An internal examination of the perimeter security equipment to ensure that the latest and most secure software loads, patches, and updates are installed; and that the perimeter equipment is

configured in the most secure manner in accordance with the guidelines and recommendations of Task Force 4, Working Group 1, Common configurations.

A review of the network owner's security practices regarding upgrade and patch management procedures and processes. This includes the use of 'staging' servers or other appropriate equipment to test patches and software upgrades before they are implemented on production networks systems.

Results of this examination must return conformance with TF4, WG1 recommendations, and show adequate procedures and processes to properly implement upgrade and patch management.

Additional details of these level 2 recommendations are to be addressed by the independent organization.

Level 3 Security Classification Recommendations:

The level 3 security classification includes all the requirements of the level 2 classification plus:

A vulnerability assessment of the 'Human Firewall'. This may include researching all publicly available information (press releases, knowledge bases, etc) of an organization to determine if any information is unintentionally provided that can be used to compromise the owners network. This may also include unannounced questioning via phone, or personal visit to support staff by unknown personnel asking questions that could reveal security weaknesses if answered, or answered improperly. For example, information that can be used to compromise a networks integrity may be as simple as revealing information on equipment types, software versions, IP addresses, maintenance window times, etc...

It is also recommended that personnel in positions of holding this type of knowledge be required to sign non-disclosure agreements detailing the types of information that is and is not allowed to be discussed, and who they are, and are not allowed to discuss this information with.

The results of the assessment of the 'Human Firewall' should not return any information that can be used to compromise the owners network.

Other details of these level 3 recommendations are to be addressed by the independent organization.

Level 4 Security Classification Recommendations:

The level 4 security classification includes all the requirements of the level 3 classification plus:

A review of the network owner's internal security practices as they relate to:

- Incident Preparedness
- Incident Investigation
- Incident Response
- Incident Alerting and Notification
- Incident Recovery

Recommendations on these practices may fall under the standards and recommendations of Task Force 4, working Group 4, Best Practices for Technical Standards.

Other details of these level 4 recommendations are to be addressed by the independent organization.

Level 5 Security Classification Recommendations:

The level 5 security classification includes all the requirements of the level 4 classification plus:

Subjecting the network being tested to professional ethical or 'whitehat' penetration attempts, or other activities designed to improperly effect the tested networks equipment or services. These attempts should be carried out at a time unknown to the network owner.

These ethical or whitehat hackers should be sanctioned and certified by the independent organization.

Also, the network owner should be able to demonstrate that they comply with the recommendations laid out by Task Force 2, Early Warning. This demonstrates that the network owner is able to share information and lessons learned with the rest of the industry for the benefit of the whole. The fact that the owning organization has reached the level 5 classification, means that this information should be accurate and reliable.

The results of these whitehat attacks should not result in an actual penetration beyond the security perimeter, and the network owner should be able to report the following minimum information to the independent organization via TF 2 recommendations within a given acceptable time frame:

1. Time of attack
2. Type of attack, penetration attempt of X device, Denial of Service Attack, etc....
3. Was attack successful in penetrating the infrastructure, or disabling equipment or services?
4. Attack target. What gear was affected?
5. Source of attack, or at least the operators best guess of the source of the attack. This may include IP and/or geographical information.
6. Are any adjustments required by the network owners procedures or processes to deal with future attacks of these types?

Other details of these level 5 recommendations are to be addressed by the independent organization.

The goal of these security classifications is to be as detailed and hard as possible while being open and generic enough that anyone owning a network, regardless of size, can achieve any level classification desired. This dictates that the details of the requirements of each classification level must be appropriate to the network being certified. For example, it is much easier for someone in a home office to demonstrate perimeter security, and security practices, than it is for a large organization with a nation-wide network infrastructure to demonstrate the same requirements.

It is recommended that the determined security certification level of each certified network be kept confidential by the independent organization, and the rating and associated information be made known only to the network owner. It will be up to the network owner to decide whether or not this information should be made public. Public knowledge of each networks certification level may lead to attempted penetration attempts of networks holding lower certification levels.

As networks receive certifications, ongoing testing must occur to maintain certification levels. Many times a network can be implemented today that is secure, but will be vulnerable tomorrow due to updated hacking techniques, network expansion or modification, and ongoing device software upgrades. It is recommended that each network be tested and re-certified at least once a year.

It will be up to industry and the independent organization to determine what levels are required of what networks before we can make the statement as a country and say, "Cyberspace is secure".

Closing Statement

As this document proposed some constructs to enable a secure network equipment deployment as an architectural model, we do not suggest being a complete reference on building a secure model. This document is to emphasize the criticality of implementing a hardware configuration that will allow for the maximum benefit of the best practices of software and management of networks.

Network security has become a pivotal role in the extremely fast moving and complex business environment which can be the most fragile place for malicious invasion. By understanding the network topologies and how they inter-relate to the software and management best practices, an understanding of a secure architectural model can be realized.

IP communication infrastructures, by their very nature are open, and can be as secure or vulnerable as is needed or mandated. When securing an existing network, the effort to increase security can be a very difficult task depending on the legacy of the network. Risk analysis has to be completed and a realistic overview of reconfiguration options should be weighed to the networks mandate for security. When architecturally designing a network from the ground up, there is far less short falls to overcome. As the network expands and more elements are introduced to the topology there must be a clear path to growth with security as a key element of future assimilation of new technology.

There is an ongoing effort from software and hardware vendors to create products that are tested for exploits and vulnerabilities, however there is nothing set in stone as of now regarding how this testing is done. Companies that implement secure architectures today must continue vulnerability testing from both the inside and outside, stay current on fixes and patches, and keep up to date on existing or upcoming threats.

Once the effort has been made and the deployment has been completed, the best-laid plan can be failed by the "Human Factor". The philosophy of security must become a cultural issue that everyone who is associated with the network must adhere to. Policies that inform staff of how they are to use the network in great detail is not an effort to be minimized. The "Human Factor" is one of the hardest to manage and should be addressed early and often.

Additional Resources

These resources may be consulted and implemented to supply, or create additional levels of security that are not covered in this recommendation.

ITU x.css

Draft Recommendations [1] T01-SG17-030910-C, draft, new ITU-T Recommendation X.css. Security Architecture for Systems Providing End-to-End Communications, ITU-T, July 2003.

ANSI T1.276

ANSI T1.276-2003, American National Standard for Telecommunications Operations, Administration, Maintenance, and Provisioning Security Requirements for the Public Telecommunications Network: A Baseline of Security Requirements for the Management Plane.

ANSI T1S1

ANSI T1S1/2003-285, American National Standard for the next generation network control and signaling plane security.

© 2004. This is an unpublished work the copyright in which vests in Marconi Communications Limited. All rights reserved. The information contained herein is confidential and the property of Marconi Communications Limited and is supplied without liability for errors or omissions. No part may be reproduced, disclosed or used except as authorised by contract or other written permission. The copyright and the foregoing restriction on reproduction and use extend to all media in which the information may be embodied. Rev. March 2004.

Marconi and Planet are trademarks or registered trademarks of the Marconi group of wireless telecommunications companies, which include MSI, Metapath Software International Limited, and Northwood Technologies, Inc. This document may contain other trademarks, trade names, or service marks or other organizations, each of which is the property of its respective owner.