



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

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

Interested in learning more?

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

**Web-Based Application Security:
Privacy Architecture for Enabling Internet Access of Patient Data**

SANS GSEC Practical 1.4

January 16, 2005

Table of Contents

Table of Contents	2
Abstract.....	3
A Quick Word About HIPAA	4
Privacy Concerns	4
Web Portals	5
Consumer Portals.....	5
Physician Portals	6
Web Applications.....	6
Privacy Architecture	6
Real-World Threats.....	8
Denial of Service (DoS)	8
Viruses	8
Quality of Service (QoS).....	9
Password Attacks (Brute Force and Dictionary)	10
Buffer Overflow	11
Open Ports.....	11
Social Engineering - and other Internal Threats	12
Man-in-the-Middle Attacks	12
Cookies.....	13
Shoulder Surfing and Key Logging.....	13
Bad Code and Other Human Errors.....	13
Security Event Logging.....	13
Encryption	14
Authentication, Authorization, and Auditing (AAA)	14
Authentication	14
Authorization.....	15
Auditing.....	15
Putting the Pieces Together	16
Measuring Success (or Failure).....	19
List of References & Internet Sources	20

ABSTRACT

Web-based applications and information portals are on the rise for healthcare organizations of all sizes. Both doctors and patients share the desire to access the necessary patient identifiable data, and seldom do they consider the security that is required to support the systems that deliver this data to them. For those that must design, implement, and support these security measures, there needs to be a solid understanding of the real security and privacy risks so that only a practical and well-structured solution emerges.

When it comes to protecting patient data and other private information, it is imperative that organizations define their security and privacy mission. This mission should be engrained in the culture of everyone that is responsible for the data that is presented to the end-users via the web interface, not just the web administrator. Consider the following sample mission statement:

Mission Statement:

Having deployed web-based services over the Internet, it is the mission of our organization to fully respect the data that we are responsible for safeguarding. We will be proactive in our security and privacy measures, work hard to anticipate relevant threats, and diligent in protecting the entire environment from unauthorized or improper access. Furthermore, we will do so with the understanding that the reputation of our organization is in the hands of those whose online data we protect and will work to provide the best balance of security, privacy and usability.

Those that have taken such a mission statement to heart will likely avoid becoming one of the many very public stories surrounding the wholesale compromise of security[], and ultimately the lost privacy of the individuals that had conveyed the trust for their very private and personal data. Of course, the public exposure of security breaches only represents a sample of the total number of compromised systems. Those that have not yet considered the consequences for having at least a moderate security model should take a few minutes right now to imagine what the headlines might look like and what your response might be; considering this to be a first high-level risk assessment.

Security, however, should not scare anyone away from deploying web-based services for your doctors, patients and other staff (both clinical and non-clinical). The face of patient care is changing to incorporate



more and more new technologies, and online interaction for healthcare patients and providers is already assured its place in this market. By addressing both security and privacy on the front-end, healthcare security professionals should be able to maintain the integrity and image of the organization.

The bottom line is this: let the application vendors market and sell the value of their online systems today, and implement the appropriate security measures that will assure that they will be available for use tomorrow. When in doubt, a risk assessment should identify the relevant threats to the patient information, compared to the overall impact for each of these threats. The results from this risk assessment will assist in determining the appropriate technologies to implement. As a starting point, this document should provide the foundation components that are necessary to understand the high-level risks for online patient data, while also outlining a privacy architecture that is geared toward both security and regulatory issues. In many respects, this document is a walk-through risk assessment that is specifically geared to healthcare security professionals and decision makers.

A QUICK WORD ABOUT HIPAA

HIPAA, the Healthcare Portability and Accountability Act^[1], is great and we all love it, right? Okay, how about HIPAA provides penalties... etc..; hey, what were we talking about anyway?

Forget HIPAA for a moment, focus on and embrace best practices for ensuring privacy and system security. What would you like to see protecting your data? Do you want to rely on government legislation to scare away those that want to hack into your systems, for whatever motivates them to do so? With or without HIPAA, it just makes sense to take all appropriate measures to protect the private data for those you serve. As you read through this document, try not to focus on how all of this matches up to HIPAA standards. Instead, focus on the real-world issues and the technology that is aimed at addressing them. Later, you can – and should – go back and see how HIPAA integrates with your plans.

PRIVACY CONCERNS

When you consider the data that you are planning to enable over the Internet, what do you consider being more important: system security or privacy? If you answered privacy, then you got it right. When it comes to patient data, privacy is what you are trying to protect; system security just relates to the controls that you put in place to ensure that what is meant to stay private, stays private. Those individuals whose patient information resides under your control seldom care about what systems you have in place; they just have the very realistic expectation that their data is only accessible by those that are properly

authorized. Your firewall could be hacked twelve times a day and not cause a stir, but that all changes once the private data is exposed.

There should be no doubt that patients have trusted their healthcare organizations to respect their privacy, including all data that resides on online systems. Unfortunately, this relationship can quickly deteriorate with even the slightest hint that their data has not been properly protected. If the morning headlines read “Patient Data Exposed”, how long will it be before all online access has to be cut off until the problem can be completely assessed and corrected? What long-term and short-term impact would this have on the overall quality of patient care? What else might come out of such an event? Will the ensuing scrutiny of your security systems and processes show that you have a deep respect for patient privacy?

It is, perhaps, somewhat ironic that hackers often exploit “trust relationships”^[iii] between servers to gain access to systems. This type of attack allows them to represent their system as being trusted by the target system, thereby gaining unauthorized access at an administrative level. When you compare this to the trust relationship that you have established between the patients and the hospital, the comparison is quite clear: an exploited trust is a major event that can destroy all credibility.

WEB PORTALS

Web portals provide a powerful front-end approach to users seeking a single point of access to perhaps multiple back-end systems. From the healthcare perspective, this is even more significant since all information is customized to suite the particular needs of a well-defined target audience: healthcare patients and care providers. From a security perspective, the fundamental difference between the two is that physician portals provide broad access to doctors seeking clinical information on all of their patients, while consumer portals are restricted so that the patients can only see his or her own data. Ultimately, the levels of security that must be implemented will vary between these two, as will be discussed throughout the remainder of this paper.

Consumer Portals

Consumer portals are no longer defined as the web page that the hospital maintains to allow the local community to learn about the hospital and its services. Although they can exist on the same physical web server, the concept of the true consumer portal is to allow patients to interact in ways that they never could before. If you are at all familiar with Internet banking, you will see this as the difference between going to the bank’s web page and using that same page to access your bank account. The bank’s web page is for information on the various services that they offer, while the web page (consumer portal) that you

go to so that you can access your bank accounts can show you the balances and other records from multiple back-end databases.

Whenever addressing the security needs of these consumer-related portals, you should check with the application vendor to see what security measures should be put in place. Similarly, there should also be certain processes recommended by the vendor or enforced within the application itself, such as having customers opt in or out of the service offering. The overall strength of the security, which is directly proportional to the controls that are put in place, really depends on what is being accessed.

Physician Portals

Physician portals are quite different from the consumer portals in that they allow the physicians to access patient identifiable data for all of their patients. Just imagine what damage could be done if a physician used his or her username along with an easily guessed password to access the portal and someone other than that physician was able to just guess their way on to the system. With the consumer portal, users could not change critical information such as their own medication. With a physician portal, the potential for serious or grave damage certainly exists and should therefore be properly addressed through the various security mechanisms. Thankfully, security solutions do exist that can address these concerns.

WEB APPLICATIONS

Web-based applications do not fall into the same category as portals. Unlike web portals, which are basically a front-end to take you to other servers or databases, web-based applications are built on a one-to-one relationship between the client application – the web browser – and the server. Web-based e-mail, such as Outlook Web Access, is one such application that healthcare organizations regularly seek to deploy for the convenience of their users.

As you seek to provide a secure method for deploying these services over the Internet, it is important to understand the different risks associated with each of the applications. The application providers should be able to provide some sound guidelines on how to safely deploy these applications over the Internet. If they say that Secure Socket Layer (SSL) encryption and a firewall is all you need, that is a good indication that you should probably ask for a more qualified opinion.

PRIVACY ARCHITECTURE

So far we have looked at the various web front-end systems that an ever-increasing number of healthcare organizations are seeking to deploy. If that was

all there was to web security, as many application providers would have you believe, you could probably just live with your SSL encryption and firewall server and all would be well. Unfortunately, just as these servers are the front-end to your data, they are also only the front-end to your security and privacy concerns.

As previously stated, the security controls are implemented to protect the privacy of the patient data. This implies that an architecture focused on privacy would serve the organization better than one that is focused on basic security, if for no other reason than it focuses on the strategic goal of protecting patient privacy, rather than the tactical goal of selecting a particular security technology. This makes more sense when you consider that data needs to be protected in more places than the web server and the database server: the data paths between the servers and the users should also be secured. And, just as important is the way that internal staff deals with the data during processing, archiving, or otherwise manipulating the data; even internal system administrators and data analysts don't need unrestricted access to all data.

Correctly authenticating users, authorizing their access to specific information and auditing their actions serves as the foundation for the privacy architecture. Not surprisingly, technology alone cannot solve all of the issues; internal staff is just as responsible for protecting the data. Fortunately, the idea of having your users become a part of the security equation now has a name: the human firewall. For an in-depth review of what makes up the human firewall methodology, refer to the web Human Firewall web site^[iv].

Overall, it is imperative that security technology is not the only focus of your efforts. Data needs to be protected at all levels, which typically requires multiple technologies to be integrated with individual efforts. The real-world threats outlined in this paper will reinforce the position that the challenges facing healthcare information security (Infosec) and privacy are both complex and varied. However, when the overall architecture is focused on privacy, as the various levels of security implemented should reflect, healthcare organizations can provide a quality mixture of security, privacy and usability.




Figure 1

When evaluating specific information security technologies to fulfill the needs of the overall privacy architecture, it is important to focus on products that not only meet the required security and privacy goals, but also do not cause an increase in user complexity. Traditionally, information security controls have caused a significant increase in user

complexity. As **Error! Reference source not found.** shows, using tools that also decrease the amount of user complexity can provide for a better user experience. Finding solutions that match this model is perhaps the greatest challenge for healthcare information technology security professionals and decision makers.

REAL-WORLD THREATS

Once you open your web-based applications to the Internet, you immediately become susceptible to a very large and ever-increasing list of information security threats. Understanding each of the major threats and addressing them within your tactical solutions will ultimately determine your privacy profile, while ignoring or otherwise dismissing any of them could end up being the entry point for attacks on your systems and data. To further illustrate the potential threats to your systems, you can view the Incidents.org web site and see an updated geographic view of what hackers are targeting. While this can – and does – change on a regular basis, it does provide some insight into the more relevant risks that are facing the Internet.

Unfortunately, there are just too many real-world threats to cover within the scope of this document. In fact, there are probably too many to cover within the binding of an eight hundred-page book. This document should, however, at least lay the baseline foundation for understanding what high-level challenges your privacy architecture will face and what methods should be used to address each of them.

Denial of Service (DoS)

Denial of Service attacks (DoS) attacks are typically very simple to achieve and are based on one thing: denying system availability by overloading resources. The big brother to this attack is the Distributed Denial of Service (DDoS), which accomplishes the same thing but with far better results since it is based on multiple hosts overloading a single system. In either case, DoS or DDoS, users and other system services are denied access to the necessary resource.

One of the more difficult challenges when dealing with DoS incidents is determining where they are coming from and whether or not they are the intentional or accidental. In most cases, simply planning for DoS attacks and implementing sufficient controls to can greatly reduce the possibility that web-based services will suffer.

Viruses

When you think about computer viruses what comes to mind? Do you see them as mere nuisances that are just part of the way the world is, or do you look

deeper, focusing on prevention and seeking to understand the nature of what the virus is trying to achieve? Do you understand each of the major groups of computer viruses: boot sector, file infecting, polymorphic, stealth and multi-partite? Are you aware of the non-virus programs that can attack your systems, such as worms, trojan horses, and logic bombs? Can you also determine what is a hoax and what is critical? Did you know that there are currently over 60,000 known viruses^[M]? While we don't address the specifics for each of these here, the fact is any of these can find their way into your organization and cause problems ranging from virtually no-impact to a complete DoS for multiple systems, including servers, workstations and network access.

When you match up any or all of these threats to your Internet strategy, what is the potential impact? A DoS situation is fairly obvious since large virus outbreaks can overwhelm system resources, while trojans, worms and other malicious scripts can perform similar attacks from the LAN portion of the network that a hacker might not have been able to achieve due to suitable firewall protection. The ability for trojans and worms alone to easily bypass even the most expensive firewall solutions is perhaps the greatest challenge to your Internet strategy.

Quality of Service (QoS)

Quality of Service (QoS) is based on emerging networking technology that can provide guaranteed bandwidth levels all the way down to a single (per-user) session. The benefit to healthcare organizations deploying web-based applications and services is significant: the user population will experience the same levels of performance during peak and off-peak times. Basically, when the network is operating under high load, specific users, or groups of users, will have a pre-determined amount of bandwidth carved out specifically for them, thereby increasing usability and customer satisfaction. Without planning for QoS issues, viruses, DoS attacks and other events that can cause unusually high network traffic, the functionality of your web-based applications and portals can easily become an escalated and time consuming matter.

Figure 2 illustrates this concept by showing multiple users trying to access systems that reside at the hospital, while each user resides at various locations on the Internet. You will notice that standard users are forced to share the hospital's bandwidth with other web services (HTTP), file transfer services (FTP), and Internet e-mail services (SMTP), while the two doctors requiring access to clinical data have been configured for a minimum of 6K/sec throughput. It is important to note that this does not fully guarantee that they will be able to achieve these levels, given that the Internet itself does not support the various vendor-specific QoS concepts. It does, however, provide a reasonable amount of control over the hospital's Internet connection so that high traffic levels both to and from the organization do not cause a DoS condition for those that require access to critical systems. The Nortel Networks Web OS platform is just one of

the emerging products that are set to deliver this solution to the healthcare market^[vii].




Figure 2

Password Attacks (Brute Force and Dictionary)

Password attacks are based on the idea of trying multiple username and password combinations in the hope that at least one account is granted access. Consider each of the following weak username and password attack scenarios:

- ❑ Several of your users choose to use passwords that are very easy to remember. They choose to go with the password of “*password*” or their own *username*. A hacker has a recent directory of all of your employees and has managed to figure out that your username naming convention is based on the first initial of the user’s first name followed by the first seven characters of the last name. So, the valid user John Doe now has a username of jdoe. The hacker takes this naming convention, applies it to each of the names in the employee directory, and uses one of the many password attack programs to mount a password attack on your systems. How many of your accounts might he be able to access by trying each of the usernames with the passwords of *username* or *password*?
- ❑ Say the hacker found 5% of the accounts were accessible using the previous attack. Delighted by this success, he or she decides to go for more, preferably some with elevated system privileges. Already having the usernames for each account worked out, and knowing that 95% of them are using something other than their *username* or *password* for their password, it is decided that a brute force or dictionary attack will be required. With these attacks, large dictionaries of common passwords and other words found in the dictionary are used with the username list to try thousands of potential username and password combinations. Since the hacker’s computer manages all this, the ability to try so many combinations is a trivial task that can be reduced to a matter of hours. Most system administrators will no-doubt perk up at this point to point out that their password policy locks out the

account after three failed attempts, which will effectively thwart such an attack. Keep in mind, however, that the accounts were locked so the by-product of the password attack, mixed with the password policy, is a very effective DoS attack: all legitimate users are now denied access.

- ❑ In a similar situation, your users have been forced to use strong passwords – those that contain at least seven characters and include numbers, letters and special characters – and many of them have registered themselves with various outside companies, using a web page for enrollment. In order to remember the username and password for each site, many of them decide to use the same username and password from work when enrolling themselves on these external sites. Anyone with access to these account databases now has the e-mail address, username and password for your users. In this situation, there is no need to use password attack programs.

Buffer Overflow

Buffer overflow attacks are one of the major strategies used by hackers today. They are relatively easy to perform, can be delivered manually or through automated means such as worms, and are highly effective at gaining access to systems on the Internet. At the basic level, they function by providing more data to a program than the programmer had made available within the system memory buffer, which is further compounded when the application does not place some specific controls on exactly what is expected and rejects everything else. So, when more data is passed to the application memory buffer than it can handle, arbitrary commands can be run on the server. The results of such attacks could include modified data, compromised data, or a denial of service.

Open Ports

There are 65,525 ports for every IP address, and each port is associated with an application or service running on the target servers. Well-known ports, such as port 80 for web sessions and port 443 for secure web sessions, are among the more popular ports that hackers look for. When the server resides on the Internet, a quick scan of the 65,525 ports on the target server(s) will reveal which ports are available – no questions asked. Now the hacker has the ports, knows – or can easily find out – what they are used for, and is ready to go deeper into your systems.

Generally speaking, the fewer ports that are available, the more secure the system is since it gives the hacker less to work with. Technically speaking, you can actually have multiple ports open and still be more secure than a server that only has one port open; it all depends on the relative security strength of each port.

Social Engineering - and other Internal Threats

If you have not yet heard that the greatest threats to network security are within the organization, this should be your wakeup call. If you are truly seeking an end-to-end privacy architecture that can protect all levels of Internet accessible patient data, the threat from internal sources cannot be overlooked. Policies, procedures and security awareness programs only scratch the surface of the human and internal technology issues that can render all of your controls useless. For example, a single wireless access point that has not been properly configured for security could allow a back door to your network so that anyone within several hundred yard of the access point could then use it to gain unrestricted access to your data center.

Kevin Mitnick, who is arguably the world's best-known computer hacker, shared the following thought while speaking to the U.S. Senate Governmental Affairs Committee:

“Companies spend millions of dollars on firewalls and secure access devices, and it's money wasted because none of these measures address the weakest link in the security chain: the people who use, administer and operate computer systems”^[viii] – Kevin Mitnick

Kevin's point here is not that these technologies are flawed, but that they can be easily bypassed by focusing on our own human nature. When those that protect the corporate data are too helpful, system security and privacy are at risk. Just like a successful con artist, those that are highly successful at hacking through social engineering all have the same basic skill: they can easily exploit someone's trusting nature to get around complex software-based security systems.

Man-in-the-Middle Attacks

Man-in-the-Middle attacks are based on the ability for a hacker to somehow place their attack computer between the users on the Internet and the server they are trying to connect to. This might sound difficult, given the fact that Internet traffic takes multiple paths, but in many cases it can actually be quite easy to accomplish, with highly effective results. The end result of such attacks could allow a hacker to perform any of the following tasks:

- ❑ Data could be collected and modified before sending it to the final destination;
- ❑ Logon information could be captured in order that the hacker could replay this information to the server at a later time;
- ❑ Encrypted SSL traffic could be routed through a separate computer in order to gain access to the data while it is unencrypted;
- ❑ All data, including both encrypted and unencrypted sessions, could be passively monitored for content.

Understanding where man-in-the-middle attacks could be successful within your privacy architecture and actively working to control these threats is an essential component, especially when dealing with SSL security on web servers.

Cookies

In the Internet world, cookies are small files that can store various forms of information on personal computers. In many respects, these can assist in provided a better user experience. On the other hand, applications that implement cookies that store private data, especially patient identifiable data, should be understood prior to deployment so that they do not become a liability.

Shoulder Surfing and Key Logging

Shoulder surfing is what happens when an individual visually monitors what someone is typing as the log on to a system, typically with the intent of using the same logon credentials at a later time. Key logging provides the same result, but does so using either software or hardware-based key logging methods to capture all typed text. And, unlike shoulder surfing, key logging methods are highly accurate and easily deployed.

When a hacker enables shoulder surfing or key logging, it can be assumed that access to critical systems based solely on static username and password combinations should not be used. When you further consider that there is no simple way to protect web-based users from such intrusions, it becomes clear that, at a minimum, username and password combinations should not be used on systems that a) contain patient-identifiable data and b) are deployed across un-trusted networks, such as the Internet. When these conditions are met, multi-factor authentication methods used only be used to authenticate users.

Bad Code and Other Human Errors

Unfortunately, we humans make mistakes. When those mistakes involve web-based applications, the potential for compromised patient data is further elevated. Common programming errors and misjudgments could include security back doors and CGI scripts that can be easily exploited, while outdated or incorrectly configured firewalls can also become security holes that further simplify the tasks for even casual hackers.

Security Event Logging

Network operating systems, servers, firewalls, intrusion detection systems and other network devices log security events, but most system administrators do not have the time or tools to actively monitor them. In the case of many healthcare organizations, where administrators are primarily available only during daytime

hours, the ability to actively monitor security events is further reduced. Without the time, tools and obligation to monitor these systems, they are perhaps never monitored. It is, therefore, not only important to log security events, but to also provide the means to monitor them on a constant basis and respond accordingly to any threats.

ENCRYPTION

The goal of using encryption to secure private information as it flows across the Internet is to ensure that transmitted information has not been corrupted, modified or otherwise viewed by anyone other than the person that initiated the connection. SSL encryption is the de facto standard for secure Internet communication.

However, aside from still being vulnerable to the majority of the attacks already listed in this white paper, it also suffers from another problem: processing overhead. Without first providing for enough system CPU horsepower while determining the long-term capacity planning, even an otherwise high-performance server can easily be brought to its knees by the overwhelming amount of mathematical processing that is incurred by the encryption algorithms. Many common servers can see significant performance degradation with only a few SSL connections.

AUTHENTICATION, AUTHORIZATION, AND AUDITING (AAA)

The ability to properly authenticate someone, authorize what resources the authenticated user can access, and then audit what has been done is the foundation for a well secured privacy architecture and is commonly referred to as a AAA solution. It is, however, generally considered good in theory but challenging in practice. In healthcare, the problem is seldom the tools or technology to enable this architecture, but the willingness for the organization to move beyond low user impact systems, such as those that only require simple username and passwords for access, to those that force multi-factor authentication on the user population and more security administration on the information services departments. Due to the extensive – although necessary – scope of these architectures, it is without question that only well documented and enforced security/privacy policies can drive AAA solutions into successful existence.

Authentication

Authentication procedures and their related tools should be portable, easy-to-use and broadly accepted by healthcare workers. Today, of the various multi-factor (what you have, such as a hardware token, and what you know, such as a PIN) authentication systems, token-based authentication is perhaps the best overall

authentication solution since it is well adapted to a highly portable workforce, especially when those users require broad Internet-based access. When patient data is being accessed over the Internet, only these highly portable multi-factor authentication systems will suffice. From a privacy perspective, the challenge is protecting the end-users from identity theft.

Authorization

Once users have provided their multi-factor authentication, they are granted broad access to the network. Now it is imperative to ensure that each authenticated user is only provided access to specific resources. In effect, this is acting as a traffic cop, pointing users only in the direction they are allowed to go and restricting all others. Finding the right mixture of authorization controls is not unlike finding the right authentication controls in that both should not overburden the users. However, since these controls can be transaction-based, user-based, or role-based, they can be matched to a broad range of applications and users.

Auditing

Auditing system access for users that have been authenticated and authorized is essential for both intranet and Internet users. For example, if patient data is exposed, only a comprehensive auditing systems can tell you if the access came from inside the firewall and under what circumstances they were allowed to access the data. It might not appear all that important to determine where the hole was after the attack, but it will go a long way in shoring up the organizational credibility if the issue is immediately understood, documented and corrected. Without the auditing components, you are left with only a few answers, and “I don’t know” or “we are looking into it” generally doesn’t go over that well. At a minimum, creating an end-to-end audit trail whenever patient records are accessed or updated should be the goal.

PUTTING THE PIECES TOGETHER

So far we have spent a fair amount of time going over many of the real world issues and best practice approaches to Internet security and privacy for healthcare data. Now, we need to put these pieces together to illustrate what a healthcare-specific privacy architecture might look like. It is important to note that this is only one possible configuration of which there might also be several very acceptable variations. Just as there is no single firewall solution that can address every security need, this proposed architecture only serves to address as many of the common challenges as possible, without focusing on each variation that might be used under different circumstances. Primarily, this solution is focused on what a generic architecture might look like that is used to support physician and consumer healthcare portals.

Figure 3 provides an overview of what a privacy architecture might look like for web-based healthcare application services.

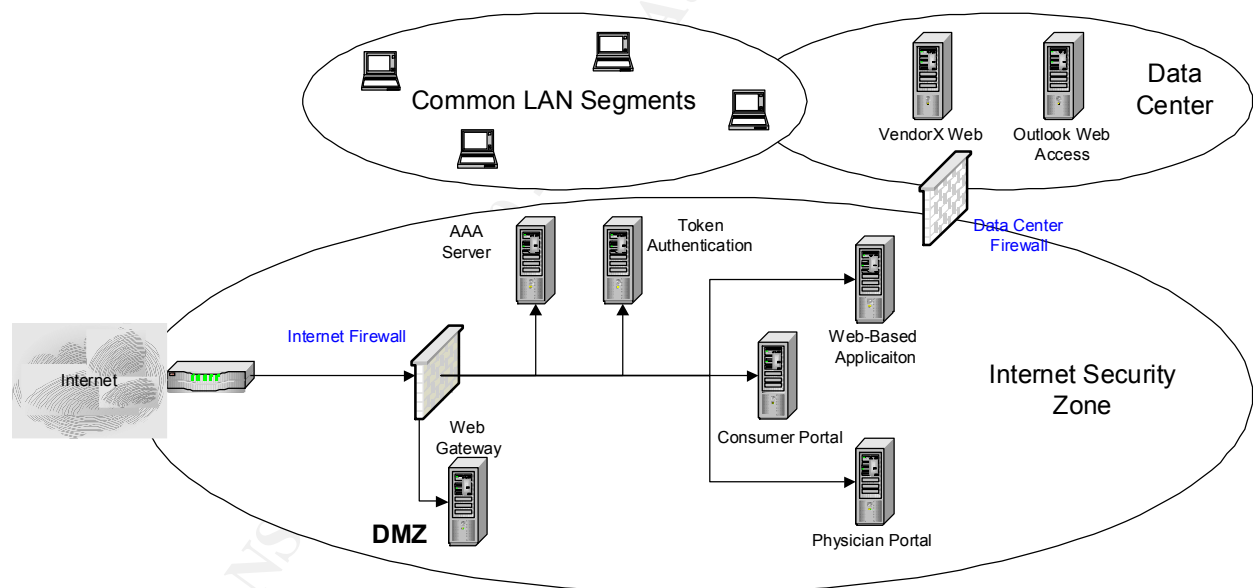


Figure 3

Perhaps the most obvious architecture change – as it relates to most environments – is that there are actually two firewalls separating the Internet accessible servers. Although these servers physically reside in the data center, they are logically separated from the rest of the network by the Data Center firewall. This configuration isolates all of the Internet accessible servers into a separate security zone, identified in Figure 3 as the Internet Security Zone.

The Web Gateway server that resides in the demilitarized zone (DMZ) of the Internet firewall is used to act as a traffic control point – or choke point - for all web traffic destined for the back-end servers. Through an interactive relationship with the Authentication, Authorization and Auditing (AAA) server, the Web Gateway server acts as the gatekeeper to each of the remaining servers in the Internet Security Zone. Without first passing through this server, users are restricted from accessing any of the servers that reside behind the Internet firewall. Since the firewall is allowed to pass the web traffic, this additional security mechanism works to ensure that application-layer attacks do not go beyond the DMZ.

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

Table 1 shows the high-level relationships between each of the components shown in Figure 3 as they relate to the security challenges that were introduced in the previous sections of this document.

Threats & Technologies	Architecture Components
Denial of Service	Internet Router & Firewalls
Viruses, Worms and Trojans	Web Gateway
Quality of Service	Internet Router, Firewalls
Password Attacks	Token Authentication
Buffer Overflow	Web Gateway
Open Ports	Firewalls
Social Engineering	Organizational Training, Policies & All Listed Components
Man-in-the-Middle Attacks	Web Gateway & Token Authentication
Web Cookies	Token Authentication, Web Gateway
Shoulder Surfing/Key Logging	Token Authentication, Organizational Training
Bad Code & Human Errors	All Listed Components
Security Event Logging	All Listed Components
Web Encryption (SSL)	Web Gateway
Authentication	Token Authentication
Authorization	AAA Server & Firewalls
Auditing	All Listed Components

Table 1

By understanding each of the security threats and privacy issues listed in the previous sections of this white paper and applying them to this privacy

architecture, a quality foundation can be implemented to address long-term privacy concerns.

MEASURING SUCCESS (OR FAILURE)

Once the privacy architecture is in place, it is important to provide regular security audits, preferably by independent information security auditors that can accurately measure your privacy architecture and rate it according to its relative security strengths and weaknesses. Should your overall privacy architecture become compromised at some later date, having a system that has been independently audited can take you a long way in overcoming the public perception that the system was compromised due to a poor design.

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

LIST OF REFERENCES & INTERNET SOURCES

- ⁱ Pallarito, Karen. "Patient Information Exposed in Health System Security Breach". Reuters Health. August 11, 2000. URL: <http://www.cancerpage.com/cancernews/cancernews1485.htm>
- ⁱⁱ Author Unknown. "Part Four: Privacy, Confidentiality & Security". HIPAA.net.com. URL: <http://www.hipaanel.com/upin4.htm>
- ⁱⁱⁱ Brunson, Drew. "Role of Perception in Information Warfare - Selected Social Aspects of Psychological Operations In Information Warfare". SANS Institute. July 16, 2000. URL: <http://rr.sans.org/infowar/perception.php>
- ^{iv} HumanFirewall.org. "Building a Human Firewall". HumanFirewall.org. © 2001. URL: <http://www.humanfirewall.org/>
- ^v Incidents.org. "Internet Storm Center". SANS Institute. © 2001. URL: <http://www.incidents.org>
- ^{vi} Avert Labs. "Virus Information Library" McAfee. © 2002. URL: <http://vil.nai.com/vil/default.asp>
- ^{vii} Nortel Networks. "Product Brief: Alteon WebOS 9.0" Nortel Networks. © 2002. URL: http://www.nortelnetworks.com/products/library/collateral/intel_int/webos.pdf
- ^{viii} Wasserman, Elizabeth. "Mitnick schools feds on hacking 101". IDG. March 3, 2000. URL: <http://www.cnn.com/2000/TECH/computing/03/03/mitnick.the.prof/mitnick.the.prof.html>