



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>

Securing e-Commerce Web Sites.

Introduction

Securing web sites, and web servers in particular, has been the focus of many security articles and conferences over the past few years. Obviously, a web site's security level is heavily influenced by the security means, which are used by, and on, the web server. It seems obvious that the key to a secure web site is the level of security achieved from security of the web server. One might have "stumbled" over a web site's database security issues if he or she was interested in DBA chores. Database security is also a well-known subject in web site security, but it is mostly documented as a standalone issue.

Building a web site is a task that involves more than one OS and more than one kind of software. Therefore, the security of the web site is achieved from the synergy of all the factors and not from the web server alone.

When I set out to write this paper, little did I know that public information regarding the "fortification" of complex web sites will be hard to come by. Only few sites publicize the internal workings of their systems, and fewer the security make-up and configuration. All this said, the question I will be trying to answer in this paper is, "How do I put all these ingredients together in order to build a secure e-Commerce web site?"

Assumptions

When building a web site we must survey the risks facing the web site from all different aspects. Not all web sites face the same "threats"; many web sites are just another collection of HTML pages in the vast cyberspace of the Internet. But, web sites conducting business, containing information (considered valuable for a malicious hacker) or holding a political view, are at higher risk than others. E-commerce web sites often hold valuable information (credit card numbers or other private, personal data) and conduct business, and are thus placed at a high-risk position.

Having recognized a web site is in the high-risk zone, we must consider the different types of security hazards:

- Denial of Service (including distributed).
- Defacement (the replacement of content on a web site, indicating it has been hacked).
- Data Theft.
- Fraud (data manipulation or actual theft).

While any of these attacks might cause revenue loss, the method of defense against each is different. Since there is no global security solution that can provide the full defensive spectrum an e-commerce web site requires, it has become extremely difficult to choose the right line of defense.

Security is a product that comes with a price tag. At first, this might be very obvious since products such as firewall and anti-virus have known pricing. However, the costs of on-going security, software-security updates, new web-site technologies etc, cannot be calculated during initial installation planning. Eventually the web site owner will have to decide what level of security will be provided, while considering the current

risks and costs involved.

Web Sites Under Attack

Web site attacks vary significantly from site to site and from hacker to hacker, and their focus has changed as well in the passing years, shifting from network level attacks to web server hacking from within the HTTP protocol itself. DoS and DDoS attacks have become a hacker-sport and can be seen in different forms; Ranging from network based DoS such as PING flooding, to full connection HTTP requests.

DoS and DDoS

When a hacker wishes to “down” a web site, all which is needed, is a computing base that can produce a larger amount of CPU-demanding activities (for example, IP floods) then the web site is capable of handling. This is true for a fully clustered web site that is connected via a T1 connection, not only for web sites with more limited resources. The attacker needs only to generate traffic that exceeds the line capabilities, and effectively the web site will no longer be available to the Internet.

Generating a large amount of traffic doesn't require having a large connection on the attacker side. The attacker may choose to use “bots”¹ or amplifiers² as the attack base. Most information regarding DoS and DDoS shows the use of network level exploits and various methods of IP based flooding. The SANS paper on the subject “*Consensus Roadmap for Defeating Distributed Denial of Service Attacks*” which can be found at http://www.sans.org/ddos_roadmap.htm, reflects these methods and the possible defense.

Recently, a new method of DDoS has been developed. Using bots to open full connections to the web site, and request an object on the web site. Using full connections compromises the identity and the origin of the attack, since the bots can be hard to trace back to their owner. These connections cannot be differentiated for all intents and purposes from ordinary requests of web browsers.

Currently there are no known defenses against DoS attacks implementing full connections (CDN³ is a partial and **extremely** expensive method that isn't feasible for most web sites). This is due to the fact that no publicly available web server or security product can fully guarantee connection originates from a “bot” and not from a legitimate connection.

Defending your web site against the more “ordinary” DoS and DDoS attacks (namely network level attacks) is a well documented art, and consists mainly of ISP cooperation with the web site owner. Most methods of defense include rate-limit of

¹ Bots are computers connected to the Internet that the attacker was able to take over fully or partially using various means. These computers then act as “robots” controlled by the attacker and can be used to initiate different types of attacks (based on the level of control gained by the malicious user). One method of taking control over PC's and turning them into bots is by spreading a dedicated virus.

² Amplifiers are computers on the Internet that have a larger Internet connection or computing capabilities and are used to amplify the attack generated by the attacker.

³ CDN – web Content Delivery Network service, provided by companies such as Akamy, Adero and Eplication.

various forms, and unwanted network traffic blocking (such as fragment blocking, UDP blocking etc).

Most of the blocks need to be performed at the ISP level, or the attacker will be able to saturate the line connecting the web site, effectively denying service to the web site.

Web Server Based Attacks

Many of the network-based attacks that create a denial of service are hard to achieve, or hold little “glory” to the attacker. This said, one must consider the fact that data theft cannot be achieved via DoS attacks. Therefore, web server attacks have become extremely popular in the past few years. Web server attacks bypass the firewall since they connect to the web site with legal network requests (i.e. TCP port 80), and are hard to trace if the web site does not employ strict log file procedures.

Web based attacks vary from web server to web server. For example: gaining control over a console on a remote MS-IIS server can be achieved using different variants of the Unicode attack, while Linux Apache server console can be controlled using a Perl test cgi attack. Other attacks and vulnerabilities through which a remote attacker can gain access to a web server while bypassing the firewall are listed in various web resources, such as www.securityfocus.com, the bugtraq mailing lists and more.

Known Web Configuration


There is no single way to install a web site that will hold all the security answers. The different ways to install and configure the different web and network components varies greatly as web sites become more complex.

A few known configurations that address the security issues are:

Configuration 1 – Basic Disjointed

A straightforward configuration, which includes the web server as a multi-homed server with one interface connected to the world and a second interface dedicated for database communications. All communications to and from the web site are maintained by the firewall while internal communications are not monitored or filtered.

Figure 1 – Basic Disjointed



Pros:

1. Simplicity and streamlining of communications.
2. Easy troubleshooting on all levels.
3. Scalability (when no n-tier⁴ architecture is needed).

4. Low cost implementation and minimum hardware.

Cons:

1. Management of the DB server requires an out-of-band⁵ communication method or web server routing.
2. Web content is distributed manually or via local scripts and applications.

Security considerations:

1. This basic configuration provides network level security (via the firewall) and DB protection (via disjointed networks).
2. The load balancer (if external hardware is used) can be used as the second level network-filtering device for extra security.
3. The use of two network cards provides low-level protection against poorly configured firewall devices (for example, fire-walking will not reveal the DB server).

This configuration provides no means of application or OS level protection. The entire security architecture is based upon the filtering devices (firewall and load balancer). If the OS hardening process is not redone frequently on a per-patch basis, the web site will be vulnerable to application and OS level hacking.

In the event that the web server is hacked the database server will be fully exposed to the hacker via the web server. This is true even if the second NIC on the web server uses a different protocol. It is recommended that a basic method of filtering be used to prevent the misuse of networking protocols.

The Compaq DISA⁶ and Microsoft DNA⁷ web site designs are similar and are basically modeled in this configuration. Both Compaq and Microsoft rely on the OS hardening process to provide the application level security and on the programmers' capability to produce secure code.

Configuration 2 – Filtered Disjointed (figure 2)

In this configuration, the addition of the filtering firewall, via the second “DMZ” on the main firewall provides an added level of security⁸. Any hacking on the web servers will provide only minimal access to the database servers. Obviously the web servers can access the database server with an appropriate ODBC connector or similar means. This configuration could potentially provide a hacker (should he be able to “own” the web server machine) limited direct data access capabilities.

Application business logic for the web site is based on a separate server to allow for

⁴ The n-tier configuration is shown in configuration 2 and is driven from the need to process business logic on a separate server.

⁵ The use of out-of-band communications means that the connection to the server is done from a different route than all other communication to and from the web site.


⁶ Found on Compaq web site at <http://www.compaq.com/solutions/internet/disa.html>

⁷ Found on the Microsoft web site at <http://www.microsoft.com/technet/treeview/default.asp?url=/technet/itsolutions/ecommerce/maintain/operate/ecomsec.asp>

⁸ This configuration can be achieved with a second firewall for improved performance. The firewall would be placed between the DB and the IIS servers (as suggested in the MS paper). It is not necessary to place the DB server in the corporate network.

easier scalability. This server may also be used for web management. Software such as MS Site Server or MS Application Server provides the content distribution, web statistics etc.

Figure 2 – Filtered Disjointed



Pros:

1. Relatively easy installation and routing configuration.
2. Easy troubleshooting for connectivity and system level events.
3. Minimal hardware.

Cons:

1. Development environment must be similar to the production web site, to allow developers to adjust application connectivity with internal servers to the filtering device used.
2. The use of one firewall as a filtering device might show a degradation in the site's performance. Should the use of extra firewalls be applied, cost and ease of installation will no longer be an advantage for this configuration.

Security considerations:

1. This configuration provides network level security (via the firewall) and DB protection (via disjointed networks). It also provides low-level application protection since core data processing is shifted from the front-end web servers to back office application servers that have no direct communications with the site's users.
2. If MS SQL is used, TCP 1433 should be used instead of named pipes. This will provide a higher level of filtering.
3. When implementing the web content distribution mechanism it is recommended not to use windows shares. FTP or MS Site Server replications are preferred.


The "Filtered Disjointed" configuration provides the administrator with the tools to filter all network-based activity on the secure side of the firewall. The main idea behind this configuration is to eliminate the ability of one server to communicate directly with the other servers. Application connectivity is allowed to provide the site functionality (web servers will be allowed communications

with MS SQL Server using TCP 1433), and no other protocol will be allowed. Although there's a performance penalty due to the extra network segments and filtering, should one of the web servers be compromised all network transactions can be logged, leaving an audit trail.

Configuration 3 – Application Protection (figure 3)

In the effort to protect the web site from application level hacking, we need to use a “higher level” filter. The filter would be used to examine the HTTP protocol, and if possible the HTTP GET, HEAD, POST, and PUT commands and parameters. This parameter should comply with RFC 2616 (<http://www.faqs.org/rfcs/rfc2616.html>) and with the restrictions of the site administrator. Such a filter can be found in some of the commercial proxy servers or in dedicated filtering products⁹. This approach apposes the Microsoft e-commerce strategy shown earlier in configuration 1, and in the e-commerce web site security, that all application level security should be driven from the DNA design and proper code writing.

Figure 3 – Application Protection



Pros:

1. High level of assurance that Internet traffic enters the various applications in the correct form and manner.
2. The use of proxy servers could improve performance, if the proxies implement a caching mechanism.

Cons:

1. Extremely hard to troubleshoot and configure.
2. High cost of hardware and initial installation.
3. The use of filter devices at the application level could cause functionality issues. This is due to the fact that the connection terminates at the proxy level and connection stickiness, session information and other client information might be misinterpreted before they reach the web servers.
4. It is imperative that the development of the application is done with full awareness to the system configuration. Not all existing web sites

⁹ A commercial filter, which acts as an application level proxy can be found at www.sanctuminc.com

can use this configuration with no application adjustments.

Security considerations:

1. This configuration provides a high level of security, both network and application level.
2. Application filtering might require the use of out-of-band management tools, since not all proxy servers can act as routers for other non-HTTP protocols.

The “Application Protection” configuration provides the administrator with multi-layer security protection. It can be used in versatile situations, and has proven itself in protecting web sites from new hazards such as Nimda and code-red (at the time of the worm release un-patched web sites using the “Application Protection” configuration would not be harmed). This protection, however, doesn’t scale easily to mega-sized e-commerce sites.

Monitoring tasks should be carefully planned. When monitoring a web site that has only one function that answers to HTTP requests in the client path, the monitor termination point is clear. In a configuration that holds many different components that receive HTTP requests it is imperative to monitor them separately and to assure that they are all up.

Summery

The job of building an e-commerce web site never stops. The web site, as the technology itself, constantly evolves. Security risks change as the site positions itself on the net, and, as the platform used by the site become obsolete.

The different web site configuration, and approaches shown in this document come to prove, that the network level protection that so many web sites have become costumed to, might not be enough. The use of advanced configurations and filtering mechanisms is currently the only way to “keep-up” with the increasing risks of conducting business on the Internet.

Companies such as Check Point that have long been identified as a packet-filtering firewall software manufacture, have developed their software to provide application filtering capabilities with the use of “Secure-Servers”¹⁰. This shows us that market leaders have identified the need for application level filtering.

Resources

- “*Web site security and Internet threats in the wild*” - <http://www.w3.org/Security/faq/>
- A description of the DISA model at Compaq’s web site. This is the theory behind Compaq’s recommended web site installation and the company’s statement on securing web sites. - <http://www.compaq.com/solutions/internet/disa.html>
- Microsoft’s web site. This site describes the internals of the Microsoft web site, one of (if not the) biggest web sites on the web - <http://www.microsoft.com/backstage/>

¹⁰ Secure Server is a term used in the FireWall-1 software literature to describe the proxy like application level filters used in the Check Point software.

- Microsoft's documentation on the proper way to install and configure an e-commerce web site - <http://www.microsoft.com/technet/treeview/default.asp?url=/technet/itsolutions/e-commerce/maintain/operate/ecomsec.asp>
- www.winntmag.com - The web site for the Windows 2000 magazine. Includes many publicly available articles concerning system and security issues.
- <http://www.faqs.org> - The greatest web site to find RFCs and other standards.
- Note – use MS TechNet as a resource for security planning. You might be surprised of what you find...

© SANS Institute 2000 - 2005, Author retains full rights.