Secure Architecture for a High transaction Web Site.

GIAC Security Essentials Certification Practical
Thomas Willingham
Version 1.4b (option 1)

Abstract

More and more Organizations are moving to using Web Server technology to handle their transaction processing needs. In many cases this involves, in part, an external facing Web Site. As an example this external Web Site maybe used as a means to receive and process customer orders for good and services. This type of Web Site usually requires several components such as application programs, application server, web server, and a database server.

External facing transaction processing Web Sites are more complicated in configuration and have more issues related to security and performance that need to be considered when planning the network and system configuration to support the site. In most cases the customer experience (i.e. fast response time, confidentiality, and availability of data) is a high priority for the organization. In most cases customer satisfaction is a major factor in the success of a project. This provides the major requirements that need to be considered when designing a Web Server farm to support the Web Site.

In most cases the data collected on the external Web Site needs to be processed or reviewed by internal applications. This need to have both external data collection and internal data processing requires careful thought when it comes to planning the system and network configuration to support the Web Site. Security and performance are the primary concerns that need to be addressed by the final solution.

This paper will look at the issues, security and performance related, that arise when planning a Web Site, which is intended for transaction processing. This paper focuses on Unix related issues but the concepts can be applied to Windows Server operating systems as well.

Server Configuration

When planning the configuration of a new Web Site we will need to consider security at multiple layers in the Server configuration and the Network Configuration. To begin with a review of the requirements for each type of server that supports the external Web Site will be needed. There are three basic types of servers needed for a transaction processing Web Site, an application server, database server and the web server. In general each server should provide one service and only one service. The Web Server should only be a Web Server and not a database server or application server. The same goes for the application server and the database server.

- Web Server - Will be responsible for handle the HTTP request from the user. This request may be for a static image or HTML document in which case the Web Server will return the data requested.
- Application Server – This will be server that will host the application programs that process the transactions for the Web Site. Applications written to run on the application server may be written in JAVA and run on the server within a Java Virtual Machine as a servlet.
- Database Server – Will be the server that will house all the data needed to process the transactions received by the Web Site. This server will also interact with the Organization's internal applications.

Attention will need to be paid to the Database Server and Application Server hosts as well as the Web Server host. Also the installation and configuration of each server should be completed on a private network that is not attached to the public LAN. This will ensure the server is not compromised before you have a chance to harden the operating system. What follows are basic general best practices for setting up and security a server.

The review will need to identify what services are needed to support each type of server. For each type of server all services that are not needed to support the function of the server will be disabled and de-installed. Each server will need to provide for remote administration. Best expected practice is to run ssh on each server. Ssh allows you to connect to the server thru an encrypted connection. This protects against network sniffing that can expose administrative passwords. As a further protection this service can be configured to only allow specific workstations to connect the host. This should be accomplished using public-key authentication as noted by CERT below.

> "**Ensure that the computer allows administration from only an authenticated host.**
> Authenticate the host in a manner that does not depend on network-resolved information such as IP addresses or DNS names, because intruders can falsify such information within packets sent to computers being administered. We recommend the use of public key authentication using a tool such as secure shell (SSH)."[2]

In the case of the Web Servers the only services that should be needed are the Web Server software and a mail transfer agent such as sendmail. In a CGI environment the Web Server would also host application programs. In that case the Web Server and Application Server are the same host. With newer JAVA based application servers, the application server can and should run on a separate host. This allows the Web Server and Application Servers to be protected separately and appropriately. All other services should be disabled and then de-installed. Examples of this would the rlogin family of programs, ftp and telnet. In the case of the database server typically only the database server software and sendmail are required. The same goes for the application server.

A discussion with the application development group will be needed to determine if the application has any special requirements for services on the host.

The basic idea here is to secure the servers as well as possible, but also to be able to identify when the server is being attacked or has been compromised so that action can be taken. In order to be able to detect an attack or compromise the System Administrator needs to know and be able to check the integrity of the operating system setup. A good way to verify the integrity of the operating system configuration is by creating a checksum database of all the operating system configuration files and binaries during the initial system setup process. This database would then be written to a read-only media such as a cdrom. Products such as tripwire and aide can be used for this purpose. Once the system is deployed a job would be setup and run on a regular basis to check the integrity of the operating system. The goal is to be able to detect a unauthorized change to the system and take action as needed.

Logging is an additional means of detecting an attack or compromise. This can be as simple as sending the syslog output to a logging server inside the LAN or as complicated as installing IDS sensors on the LAN segments for each of the servers. A good solution would be to do both. Send the log output from the servers to a central logging server. This server will have programs installed to monitor the logs for activity that is outside of normal activity for the application. The log server will need to make sure that the logs are protected from modification. If the server or the log server identifies a message that is not normal, these servers can print the message to a printer. This will preserve the logged message in the case of a server compromise. This data can then be analyzed to help determine how the server was compromised. The logging server can then send alerts to the System Administrator Groups. In addition a series of IDS sensors could be set up on each LAN segment. The IDS sensors would monitor the Web Server, Application Server and the Database Server. They should be set up in such a way to monitor for both external attacks and internal attacks.

Network Infrastructure

As a transaction processing based system the Web Site now becomes an important resource within the organization. It becomes necessary to protect both the Web Site from the public Internet and the internal network from the Web Site. First and for most the organization should follow the advice of “Place the Web server on a subnet isolated from public and internal networks”[2]. There are a couple of approaches to protecting the Web Site itself.

External DMZ segment

One approach is to place the Web Servers supporting the Web Site on a LAN segment that is outside the corporate firewall. This LAN segment would be directly accessible from the public Internet. The only device between this segment and the Internet would be a router. Routers can be configured to do some filtering of network traffic. But they are not well suited for this task.

This approach will make protecting the Web Site from intrusion more difficult as the host is directly accessible from the public Internet. There will be little protection for the Web Server against Denial of Service attacks or attempts at intrusion, aside from the protection offered via the host operating system. The host operating system should configured as outlined above and have only the services running that are needed to support the Web Server. Basically what remains to be done is to be able to detect an attack that is in progress, and have a plan to deal with the attack.

A key to detection is logging from the system and web server software. Logs for the server should be sent to a logging server as outlined in the above host configuration suggestions. On the logging server packages could be run to monitor the logs for intrusion attempts. One example is the open source package swatch. An administrator would need to monitor the output from these Intrusion Detection packages and act if an attack has or is occurring against the Web Site. A key issue with the setup of a logging server, in the DMZ scenario, is how to get the log output to the server in the first place. One method would be to open a hole through the corporate firewall to allow the log output to be transmitted to an internal log server. The firewall would need to verify that the data being received is from the Web Server and the Web Server only. There are several ways to accomplish this verification; however most if not all would still be vulnerable to disruption by a potential attacker due to the fact that the Web Server is outside the firewall and is not protected. One way to set up the logging server that is safer than most is to have a second LAN segment. This segment would be a private network with no external connections. In this case the Web Server would be dual-homed and the logging server would be connected to this private network. However an administrator would need access to the console to be able to monitor the logging server (i.e. possible need physical access to server). This would make intrusion detection more difficult and would slow down response time in the case of an intrusion.

Assuming that the Application Server would be more critical to protect, the Application Server would be placed inside the firewall. In this case a hole would need to be opened through the firewall to allow the Web Server to talk to the Application Server. This hole should be configured to allow only the Web Server to connect to the application server. The problem here is, how do you verify that the connection attempt is from the Web Server? One may be inclined to use the IP address of the Web Server for this verification. However the IP address can be spoofed or hijacked. There will only be a router between the Web Server and the external Internet. This router will be able to filter out some types of attacks but not all. A better way may be to set up a VPN tunnel between the Web Server and the firewall. This tunnel would require a key-based authentication of the server. In either case the Application Server would be on the internal network. There would be little protection of the internal network in the case that the Application Server is attacked. An attack on the Application Server can be launched through the Web Server by trying to exploit bugs in the application programs running on the server. We will discuss this scenario later in this document.

Given the criticality of the server and the problems of a timely response to intrusions this is not a recommended network configuration.

Firewall DMZ segment

Another approach would be a network configuration with multiple LAN segments and firewalls. Within this approach the Web Server and the Application Server would exist on a LAN segment that is separate from both the external Internet and the internal LAN segments. The Web Server/Application Server LAN segment would be behind the outermost corporate firewall. The firewall should be configured to allow only HTTP traffic to the Web Server and no traffic directly to the Application Server. Then a second firewall would be set up protecting the Internal LAN from both the Web Server/Application Server and the external Internet.

However, many of the same issues raised in the discussion “External DMZ segment” would still need to be addressed. Each host would need to be configured and hardened based on the recommendations within the “Server Configuration” section above. These would include the Web Server, Application Server, and the Database Server. Host-based Intrusion Detection should be installed and configured. Also, a Network Intrusion Detection system should possibly be set up on the LAN segment.

Setting up the host based on these guidelines will give an added layer of protection in the event that an attacker breaches the corporate firewall. Also, as outlined above, the host should be set up to send log output to a logging server. In this case the log output from both firewalls would also be sent to a logging server. Intrusion detection then would rely on both sets of logs. The firewall logs would show that the firewall itself was being attacked and the host logs would show if an attack was attempted on the Web Server or Application Server. The logging server in this case could be on the inside of the internal firewall. But the logging server should also be configured as outlined in the above guidelines to protect from internal attacks. The logging server could still run intrusion detection packages such as swatch. Packages like swatch monitor the logs for suspicious activity and can send an email alert or report giving the system administrator a more timely notice that a problem is occurring.

Within this network layout the database server will reside behind the Internal Corporate Firewall. In order to function and integrate with internal system the Application Server will need access to the corporate database server. This will necessitate opening the database port through the Internal Firewall to allow the Application Server to access the corporate database server.

With this approach the network has multiple layers of defense against intrusion. Now an intrusion can be defended against at the external firewall, the Web Server/Application Server hosts and the internal firewall. The attacker would need to breach the defenses at each of these points to be successful. At each point there is the likelihood that the attack would be detected and action can be taken. Making the attack more difficult to complete.

Web Server Clustering

At this point the Organization will have deployed a Web Site that is based on a relatively security infrastructure. This Web Site may contain a total of three servers. Each server will have facilities to log activities and to help monitor for intrusion attempts. However, as the popularity of the Web Site grows, the site will need to be scaled to support the additional load. In this case the Organization can take one of two approaches to scaling the Web Site. One approach would be to replace the old hardware with new hardware that has a larger overall configuration (i.e. more CPUs and/or more memory). Eventually, with this approach, the Organization may reach a point where a single server will not be able to handle the load required to support the Web Site. The second approach would be to incrementally scale by adding additional servers, where needed, to handle the processing requirements. This approach will require some sort of clustering solution. Clustering can be accomplished at the Web Server layer and/or at the Application Server layer. In either case we are now adding more servers to the Web Site. This approach provides scalability beyond what any one host can be configured to support. Also this is a less costly approach to scaling the Web Site overall.

As more servers are added to the support the Web Site, it becomes more complex to secure the Web Site. For one thing, with the additional servers, there may be additional services required to support the Web Site. For instance if the Application allows the user to upload files to the Web Site, there would need to be a central file servers to store the files. This would require a file sharing service, such as NFS, to be enabled on the Application Servers. The use of Web Server clustering can help to mitigate the risks associated with running these additional services and could provide some additional protection as well.

Basically web server clustering offers the ability to create what the end user sees as a single Web Site that spans across multiple physical hosts. Providing client-side transparency. Most solutions on the market also provide server-side transparency, which means that the servers can use any Web Server software package to run the Web Site. One host, within the cluster, will act as “proxy for incoming connections”[7]. For the sake of discussion this host will be referred to as the dispatcher. The dispatcher will be configured to accept connections on the Web Site’s external IP address, which is referred to as the Virtual IP Address (i.e. VIP). The dispatcher will receive all incoming requests and than using load balancing algorithms will decide which server to send the request. The server receiving the request will process the request and send the client a response.

Given that the dispatcher receives all incoming connection requests that puts the dispatcher in a unique position. The dispatcher can act as a filter to help provide some protection against Denial of Service attacks. When the dispatcher receives a new connection request, the dispatcher creates an entry in a table that contains key information about the connection. When the dispatcher receives packets that are not new connection requests, the dispatcher can compare the packet against its internal table. This enables the dispatcher to drop incoming packet fragments that are not part of an existing session. This can help protect against packet fragmentation attacks on the Web Server. Also some solutions, such as Resonate, use delayed-binding. The full handshake for the TCP sessions is handled by the dispatcher's device driver, which uses dynamically allocated tables to manage the sessions, which helps protect the site against SYN-Flood attacks as illustrated by the below quote. This quote is talking about the product Resonate, which is commercial software based Web Clustering solution.

> "Since all VIP traffic first contacts the Scheduler, and the Scheduler is doing delayed binding, all SYN attacks are isolated on the Scheduler. The RXP device driver on the Scheduler is then able to protect the site against SYN attacks because it is able to dynamically allocate additional memory to use for holding the SYN packets. The RXP device driver uses less memory to store each SYN packet than typical TCP stacks. These design features, along with the device driver's ability to actively flush SYNs that have been unacknowledged for too long, help protect a Central Dispatch site from SYN flood attacks. As long as the SYN attack is not filling up the site's network connection to the Internet, legitimate traffic can still reach the site and be scheduled to servers during an active SYN attack."[1]

As noted previously, the application may be more that just a data entry application. The nature of the application may be such that user will be able to do more that just enter an order for goods. The user may be able to upload complete documents to the Web Site for internal processing by the Organization. These uploaded files would need to be placed in a central storage location that all servers in the cluster can see. One way to accomplish this is to use the NFS file sharing service. The administrator will need to secure the NFS service as well as possible. However the NFS security depends solely on host based trust relationships. The administrator can export the NFS shares as read-only, however in the example above it is required that the NFS shared be exported read-write. There are two ways an administrator can define the trust for NFS, one would be using the DNS host name and the other would be using the IP address of the trusted host. In the former, it is possible that the DNS cache can be poisoned allow an attacker's host to connect to the NFS server. In the latter, it is possible that an attacker could hijack the IP address of the trusted host. In both cases once the NFS share is mounted by the attacker, the attacker would be able to read, write, and delete files on the NFS server.

With Server Clustering technology it is possible to place the clustered servers on a private (i.e. non-routed) network. This would leave the Virtual IP address as the only external facing IP address. “This type of configuration provides pseudo NAT functionality and enhanced security for the site”[1]. All requests destined for the Web Server/Application Server farm would have to go through the Web Clustering interface. This allows the Administrator to have fine-grained control over what services are published to the outside networks, both external Internet and internal LAN. Based on the “Server Configuration” discussion above the Administrator should configure the Clustering interface to publish only the services required to support the Web Server and Application Server. This usually will be HTTP and SSH. With the servers being placed on a private network that does not publish the NFS services outside of this network the risks associated with running this service are greatly reduced.

This private network configuration will make outbound network connects more difficult. Some examples of outbound network connections required will be logging, e-mail delivery and database connections. One solution to these outbound connection requirements could be to install one more firewall. This firewall would be connected to the private network and the internal corporate LAN. The firewall should be configured to allow only outbound connections from the Application Servers and to support only logging, e-mail delivery, and database connections. This increases the complexity of the network layout but provides an additional layer of security protecting the application servers from the internal LAN as well as protecting the internal LAN from the application servers.

Web Site clustering helps provide the performance and security Web Sites require in today’s Internet. Clustering allows the Web Site to be scaled incrementally as demand on the site increases. It can also provide a layer of abstraction between the physical server layout and the logical. Improving protection of the individual servers from DOS attacks and intrusion attempts.

Web Application

Most of the discussion to this point has been about addressing security from a network and server prospective. We have looked at creating multiple layers to protect the critical services within the Organization. But at this point the Organization now has a web presence that may include static HTML as well as dynamic content supplied by application programs. Such an application becomes a network-based service and as such becomes vulnerable to network-based attacks. All the defenses put in place so far allow HTTP traffic to pass through to the application with minimal verification. Exposing the application directly to the external Internet. The application can be vulnerable to the same type of attacks used to attack other network services, such as sendmail. Attacks can take the form of buffer overflow, SQL injection and Denial of Service. It is important to consider security within the application programs during the design, Q&A testing and deployment phase of the application life cycle.

As a network service the application should take some basic precautions. In general the application should minimize trust of the external clients. All input coming from the client browser should be verified before the information is used or written to the database. This verification will need to be done against all input fields returned from the client browser. This includes, in the case of CGI Form processing, hidden fields that are used to maintain state as well as user data entry fields. It is possible for an attacker to save the HTML produced from the CGI program and then modify the hidden fields before sending it back to the server. It would make things easier on the programmer if the state of the session on the server.

One issue to watch out for is if the input data is larger than the internal buffers that are used to process the input. If a check is not made than an attacker can case a buffer overflow within the application depending on the language used to write the application. Some languages, such as perl and JAVA, will allocate memory dynamically based on the size of the data placed into a variable in which case this would not be an issue. Other languages, such as C, do not dynamically allocate memory. In which case an attacker could fill up a memory buffer to the point of overflow. This overflow may allow the attacker to insert code into the execution stack of the running process on the server. At which point the attacker has access to the Organizations server and is executing programs of his/her choosing.

Another issue the programmer should check for, when the data is returned from the client browser, is that the data may contain escape characters. These could be SQL escape characters and/or shell escape characters. If the user supplied data is passed directly to the databases SQL interpreter. It may be possible for an attacker to trick the SQL interpreter into terminating the SQL the programmer wrote and executing SQL the attacker created. This may allow the attacker to retrieve corporate sensitive information directly from the database server.

It is therefore important that security be a design requirement[4]. As a design requirement the developers will need to consider security issues during the development phase of the application. Making it much easier to identify and to resolve security issues before the application is deployed. Also as a design requirement, the requirement would need to be reviewed during the Q&A process to verify that the requirement has been meet. In order for the developer to address security issues and the Q&A group to identify when there is a problem the Organization may need to train each group secure programmer practices. In addition a periodic review of the production application could be conducted to check for security vulnerabilities within the application. There are companies that will come in to conduct this audit of the production system.

Conclusion

When it comes to planning the deployment of an external Web Site security needs to be considered at multiple layers. From the external firewall to the applications that execute to support the transaction processing requirements of the Web Site. Careful thought needs to be given to each component that makes up the Web Site and it's placement on the network. And with each decision there are risks and trade-offs. Some of the requirements that the Organization has for the Web Site may mean running services on the hosts that have more risk associated with them. By understanding the risks, there may be ways to lessen them.

It is also important that security be considered within the application programs as well. Application programs can provide a window into an Organization's internal network. The programs may be vulnerable to SQL Injection attacks and/or buffer overflow attacks. In the case of a SQL Injection attack the attacker will be able to run random SQL against your corporate database. This database may contain data that may include personal information about your customers. This information could then be used to either steal your customer's identity or buy products using your customers credit card information. The Organization would create ill will with its customers if it allows this to happen, and might find itself losing its customer base.

Security is increasingly important on the Internet today. Attackers are creating ever more sophisticated attacks. An Organization needs to have security in grained in ever aspect of the operation of a Web Site. Security needs to be a continual process of monitoring and evaluating new threats. This paper outlines a start when it comes to setting up the network and servers supporting a Web Site but the security needs to continue to evolve in response to new threats.

References

[1] Wilder "Resonate Central Dispatch™: An In-Depth Technical White Paper" URL: http://www.resonate.com/solutions/pdf/cd_techwhitepaper.pdf (October 4th, 2001)
[2] Julia Allen "The CERT® Guide To System and Network Security Practices" URL: http://www.cert.org/security-improvement/ (May 30, 2001)
[3] Lincoln D. Stein & John N. Stewart "The World Wide Web Security FAQ" http://www.w3.org/Security/Faq/index.html (February 4, 2002)
[4] David Wong "Building Secure Systems" http://online.securityfocus.com/infocus/1596 (June 20, 2002)
[5] Charl van der Walt "Assessing Internet Security Risk: What is Risk Assessment?" http://online.securityfocus.com/infocus/1631 (June 11, 2002)
[6] Razvan Peteanu "Best Practices for Secure Development" http://www.owasp.org/whitepapers/best_prac_for_sec_dev4.pdf (Oct 2001)
[7] Trevor Schroeder, Steve Goddard, and Byrav Ramamurthy "Scalable Web Server Clustering Technologies" "http://netlab18.cis.nctu.edu.tw/html/paper/2002_01_03/ScalablewebSweverClusteringTechnologies.pdf" (January 2002)
[8] Simson Garfinkel, and Gene Spafford "Practical Unix Security", O'Reilly & Associates, Inc., June 1994