



INTERNET
SECURITY
SYSTEMS

**Security Architecture and Incident
Management for E-business)**

By Marc S. Sokol, with contributions from David A. Curry

Last Updated: 5/17/2000

TABLE OF CONTENTS

1.	BACKGROUND	3
1.1	SECURITY AS AN E-BUSINESS ENABLER	3
1.2	INFORMATION SECURITY AS DEFINED BY BRITISH STANDARD 7799 (BS7799)	3
1.3	THE PRINCIPLE OF LEAST PRIVILEGE	3
2.	COMPONENTS OF A SECURE INTERNET/EXTRANET ARCHITECTURE	4
2.1	NETWORK SEGMENTATION	4
2.2	FIREWALLS	5
2.3	CLASSIFICATION	7
2.4	AUTHENTICATION	7
2.4	ENCRYPTION	7
2.5	INTRUSION DETECTION SYSTEMS (IDS)	7
2.5.1	NETWORK-BASED IDS DEPLOYMENT CONSIDERATIONS	8
2.5.2	HOST-BASED IDS DEPLOYMENT CONSIDERATIONS	9
2.6	HOST-BASED SECURITY	9
3.	DEPLOYING THE ARCHITECTURE	9
3.1	FIREWALL	10
3.2	PUBLIC DMZ	12
3.3	PRIVATE DMZ	13
3.4	ISOLATED IDS SEGMENT	13
3.5	INTERNAL NETWORK	14
4.	INCIDENT MANAGEMENT AND RESPONSE	14
5.	CONCLUSION	17
	ADDENDUM A	18

1. **Background**

1.1 **Security as an e-business enabler**

Today, many organizations are realizing that in order to compete in a global market they must migrate key business processes to the Internet. This idea, commonly referred to as e-business, is a major transformation for many businesses, but a necessary one. Organizations are leveraging Internet technologies to broaden their market share globally, to enter into new or extended fields of business, to increase employee productivity, and to build and merge partnerships and joint ventures regardless of location. However, as organizations begin to exploit the benefit of the Internet and web technologies, they are quickly learning that there are inherent risks involved in connecting their organization's networks to the Internet. Specifically, the process of exposing valuable corporate systems and data to a wider audience significantly increases the risk of attack. As this transformation occurs, an organization's dependence on security, availability, and manageability significantly increases. Hence, security not only plays the role of protector but also of e-business enabler.

1.2 **Information Security as Defined by British Standard 7799 (BS7799)**

BS7799 is a compilation of information security best practices that was developed as a result of industry, government, and commerce demand for a common framework to enable companies to develop, implement, and measure effective security management practices and to provide confidence in inter-company trading. It is based on the best current information security practices of leading British and international businesses and has met with international acclaim. Discussions are ongoing seeking to adopt BS7799 as an International (ISO) standard. According to BS7799, information security is characterized by the preservation of the following three components:

- **Confidentiality:** ensuring that information is accessible to only those authorized to have access;
- **Integrity:** safeguarding the accuracy and completeness of information and processing methods; and
- **Availability:** ensuring that authorized users have access to information and associated assets when required.

As a result, information security protects information from a wide variety of threats in order to ensure business continuity, minimize business damage, and maximize return on investments and business opportunities.

1.3 **The Principle of Least Privilege**

The preservation of confidentiality, integrity, and availability is achieved through the implementation of a number of security practices that rely on the fundamental information security principle known as the "rule of least privilege." The rule of least privilege states that any

object (user, administrator, program, system, etc.) should have only the privileges it needs to perform its assigned tasks. BS7799 states specifically that the “allocation and use of privileges (any feature or facility of a multi-user information system that enables the user to override system or application controls) should be restricted and controlled.” Least privilege is an important basis for limiting an organization’s exposure to attacks and the damage they can cause. However, it is this fundamental principle that attackers try to circumvent – their goal is to gain privileged access to system resources. Privileged access usually means gaining what is known as “root” or “administrator” access to the system in question. In general terms, to gain such access, the attacker exploits an application or operating system vulnerability, or a poorly configured system, and obtains full control of and unrestricted access to the system. The attacker can then perform any task on the system, including compromising the confidentiality, integrity, and availability of the data that resides on that system. In addition, he can attack other systems that trust the compromised system and introduce malicious programs that can spread throughout the organization.

As new vulnerabilities are identified and system and network configurations change, the risk of attack can never be completely eliminated. Organizations must implement a secure Internet/extranet architecture that incorporates information security best practices and develop an incident management and response process. Without this architecture and process, organizations cannot effectively manage the risks associated with e-business.

2. Components of a Secure Internet/Extranet Architecture

A secure Internet/extranet Architecture includes the following components:

- Network Segmentation
- Firewalls
- Authentication
- Encryption
- Intrusion Detection
- Host-based Security

The controls provided by this architecture offer a “force-multiplier” effect – each control supports the others, thereby strengthening the overall security architecture.

2.1 Network Segmentation

In most cases, organizations should deploy at least one “Demilitarized Zone” (DMZ) network. A DMZ network is situated between the hostile outside network (e.g., the Internet or an Extranet) and an organization’s internal network. The DMZ network contains publicly accessible systems, such as web servers, mail servers, and name servers. The DMZ network is protected from the outside network by a firewall, and is monitored with network intrusion detection technology. The systems connected to the DMZ network should also be monitored at the host level for security vulnerabilities and unauthorized use.

Due to the complexity of various e-business applications, more than one DMZ network may be required to provide different degrees of access. Each set of systems or applications for which access requirements differ should be attached to its own DMZ network. Firewalls and authentication technologies are used to control access between these networks. Further, switches are used in place of hubs or other shared network technologies to limit the ability of hosts to receive network traffic that is not specifically directed to them.

As an example, e-commerce applications that use databases should use two or more DMZ networks. The application servers accessed by customers should reside on one DMZ network to which general public access is granted. The database servers, which should not be accessible to the general public except through the e-commerce application, are connected to a second DMZ network. Access to this second DMZ is restricted to the application servers on the publicly accessible DMZ. Furthermore, access to the second DMZ is restricted to protocols needed for database access. The benefit of this design is the scope of an attack is limited to the systems on the publicly accessible DMZ. Although database access is available from these systems, the ability to compromise the database systems is limited by the firewall. The internal network is also protected, because the firewall does not allow any connections to be initiated from an external network (e.g., the aforementioned DMZ networks) to the internal network. This concept will be discussed later in this document.

2.2 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. An enhancement to basic packet filtering is stateful inspection. This technology keeps tracks of “conversations” through the firewall to ensure that the only traffic allowed from the 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).
- 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.
- 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.

2.3 Classification

Data, systems, and networks must be classified in terms of confidentiality, integrity, availability, and criticality to the organization. As a result, organizations should document policies and procedures to identify a methodology for assigning this data classification.

2.4 Authentication

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.

2.4 Encryption

Encryption is an essential part of e-business 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.

As an example, consider credit card data. When ordering products over the Internet, credit card data is routinely encrypted when it is transmitted. However, there have been several cases in which credit card numbers have been stolen by an attacker from merchants who, once they received this data, stored it in unencrypted form.

2.5 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 protect all systems on an entire network segment. This capability generally makes network-based IDS easier and less expensive to deploy. However, network-based IDS is limited because it cannot "see" what is happening on individual hosts. For this reason, a complete intrusion detection

implementation will make use of both network-based and host-based solutions.

2.5.1 Network-Based IDS Deployment Considerations

An intrusion detection system should be deployed on each DMZ network, 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. 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.

2.5.2 Host-Based IDS Deployment Considerations

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.

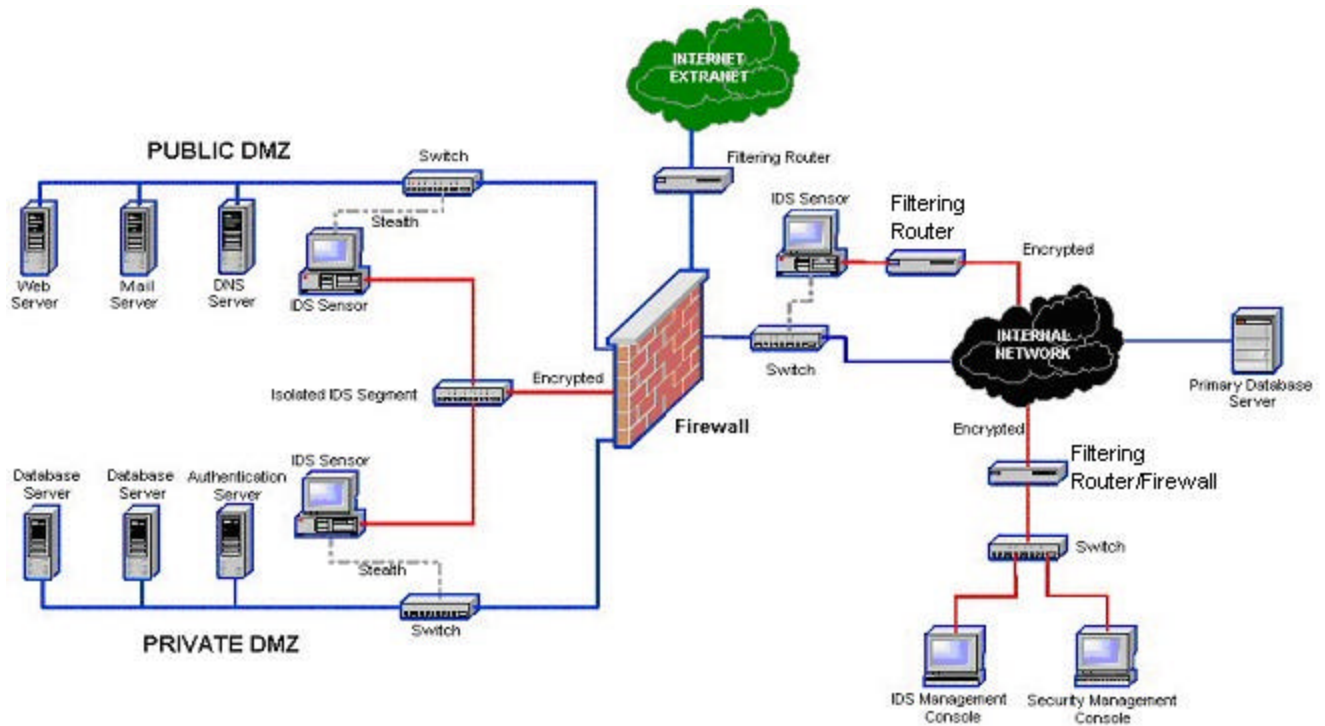
2.6 Host-based Security

Host-based security is a critical aspect of an organization's security posture. Even though hosts are protected by firewalls, authentication mechanisms, and intrusion detection, it is still possible for an attack to occur. It is therefore important that security controls also be present on the host itself.

An acceptable baseline security level must be defined, implemented, and monitored to identify and address risks and maintain policy compliance. Host-based vulnerability assessment and configuration scanning should be performed on all production systems on a regular basis. This approach ensures that operating systems and applications (e.g., web applications, database applications, etc.) are up-to-date and properly configured. Vulnerability assessment and configuration scanning can also help ensure that user accounts are properly configured, authentication requirements are enforced, file integrity is validated, object permissions and privileges are controlled, and logging/auditing is being properly performed. In addition, the system clocks of all hosts (including the firewall and intrusion detection systems) should be synchronized to a common time to allow for correlation of log data from multiple systems. These procedures ensure the accuracy and auditability of the data that has been recorded. All system audit trails and activity trails, as well as data from the vulnerability assessment scanning should be securely transferred (encrypted and authenticated) and stored in a secure central location rather than on the individual hosts. This prevents an attacker who has gained access to the system from altering or destroying the evidence of his actions.

3. Deploying the Architecture

The components described in the previous section are deployed as shown in the diagram below.



3.1 Firewall

The above diagram illustrates a single firewall with five network interfaces. It is also possible to use multiple firewalls with fewer network interfaces for performance or redundancy reasons. However, in most instances, a single firewall should provide adequate performance.

A packet-filtering router is installed between the Internet/extranet and the firewall. This router provides the first line of defense by blocking unnecessary inbound traffic as required by the following Internet standards:

- **RFC 1858:** Security Considerations for IP Fragment Filtering
- **RFC 2267:** Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing
- **RFC 2644:** Changing the Default for Directed Broadcasts in Routers

Connections from hosts on the Internet/extranet are controlled as follows:

- Hosts on the Internet/extranet 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/e-commerce 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.

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

- Hosts on the Internet/extranet 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/e-commerce 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.
- 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 IDS segment are controlled as follows:

- No connections may be initiated from the Internet/extranet, 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., e-mail/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/extranet, 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."

3.2 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 best practices approach to doing this uses “staging servers” on the internal network. These servers are configured in exactly the same way as the servers on the Public DMZ. Any changes to the servers are first tested and security validated on the staging servers before being deployed on the Public DMZ (this includes the initial system image).

3.3 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.

Although the systems on the Private DMZ are not directly accessible to the public, they are still less protected than systems on the internal network. Therefore, like the systems on the Public DMZ, these systems should also be updated from staging servers on the internal network. Any changes to the servers are first tested and security validated on the staging servers before being deployed on the Private DMZ (this includes the initial system image).

3.4 Isolated IDS Segment

The isolated segment provides for secure and controlled communication between intrusion detection systems and its associated management console. Essentially, this is an additional external network that is connected and controlled by a firewall.

Each 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 (“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 firewall 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.

3.5 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.

4. Incident Management and Response

The architecture described above is a best-practices approach to reducing the risk of penetration and limiting the damage that successful penetrations can cause. However, no architecture can prevent all intrusions. Hence, a risk always remains. It is this risk that organizations must be prepared to address through an incident management process. BS7799 states, “Incident management responsibilities and procedures should be established to ensure a quick, effective, and orderly response to security incidents.”

Organizations must develop processes and procedures to manage and respond to security incidents that occur. Security incidents can occur at any time, can cause significant outages, damage, and financial loss, and are frequently extremely complex. For these reasons, it is essential that the

incident response process be developed and integrated with the deployment of the aforementioned architecture.

An incident management and response plan must address the following issues:

- **Incident Preparedness**
 - Definition of roles and responsibilities for personnel handling incidents – A computer security incident response team must be developed. This team includes representatives from at least the following departments: IT Security, System and Network Administration, Disaster Recovery, Legal, Human Resources, Public Relations, and Corporate Security.
 - Education and training for incident management and response – Incident handling requires specialized skills. Careful attention to the manner in which an investigation is performed is essential.
 - Scenario testing and validation (“fire drills”) – The incident response process should be tested regularly.
- **Alerting**
 - Define possible sources of alerts – Alerts can come not only from technology sources such as firewalls and intrusion detection, but also from sources such as human resources (e.g., disgruntled or terminated employees), end-users (both internal and external), news agencies, attackers, and Internet service providers (ISPs).
- **Report and Notification**
 - Define methods for receiving and logging alerts – Visual alerting consoles, e-mail, pager, phone, personnel, and other resources can receive Alerts.
 - Define methods for communicating alerts – To ensure that incidents are reported in a timely manner, an incident reporting and communications process should be defined and disseminated. This will ensure that incidents are identified and that appropriate management channels are used.
- **Preliminary Investigation**
 - Determine appropriate preliminary courses of action – As a result of the callout process, roles and responsibilities are applied to managing the incident.
 - Determine whether the incident is of an emergency nature – Examination of activity logs, interviews with users and/or administrative staff, and examination of policies can help determine whether an incident is “real” and whether it is pervasive in nature. Specially trained personnel must perform this investigation to avoid contamination or destruction of evidence.
- **Decision and Resource Allocation**
 - Emergency Declaration – As a result of the preliminary investigation, a state of emergency can be declared. Staff resources can be assigned and budgetary resources can be allocated.
 - Define Coordinator – A single person is designated to coordinate response to the incident.
 - Determine whether legal action is planned – Methods of response are dictated by the requirements of any future legal action. If legal action is planned, additional procedures must be followed to

preserve evidence for admissibility in court. The decision to initiate legal action cannot be decided “after-the-fact.”

- **Response**
 - In-depth Investigation – The preliminary investigation is expanded as necessary. This may include the involvement of additional personnel, other companies, and law enforcement agencies.
 - Containment – To prevent the incident from spreading, compromised systems may be isolated and/or disabled.
 - Legal – Depending on the incident, it may be necessary to involve corporate legal council, public/media relations, human resources, etc.
 - Communication – If the incident impacts users, they may need to be informed that an incident has occurred.
- **Recovery**
 - Eradication – Remnants of the incident (attacker toolkits, Trojan Horse programs, viruses, etc.) must be removed from the system(s) affected by the incident.
 - Operations Restoration – Systems must be rebuilt, recovered, or replaced.
 - Mitigate Reoccurrence Risk – System vulnerabilities and inadequate controls exploited by the attacker must be addressed.
- **Lessons Learned**
 - Documentation – The incident, its causes, and its effects must be documented. Actions taken during the response must be documented and analyzed for successes and failures.
 - Update Process – The incident response process should be reviewed and updated as necessary.
 - Financial Impact Analysis – The costs associated with the incident should be determined. This may impact future budget allocations for information security.
 - Staff Needs – Responding to an incident may reveal the need for additional staff or the need to better train existing staff.
 - Budget Needs – The financial impact analysis may reveal that spending additional money on preventive measures is justified.
 - Quality in Information Security – The overall organizational approach to information security should be reviewed and enhanced as necessary to ensure that adequate controls are in place and monitored.

Just as the development and execution of disaster recovery plans often requires expert assistance from firms specializing in the field, so does the development of an incident response plan. Furthermore, many organizations outsource their disaster recovery operations. In a similar manner, incident management and response can be outsourced, allowing organizations to obtain assistance from entire teams of subject-matter experts at a lower cost than duplicating the entire team in-house. This gives the organization the full benefit of the experience of reacting to and resolving break-ins at other companies. Organizations outsourcing this “extended” incident response capability must still maintain their core incident response team.

5. Conclusion

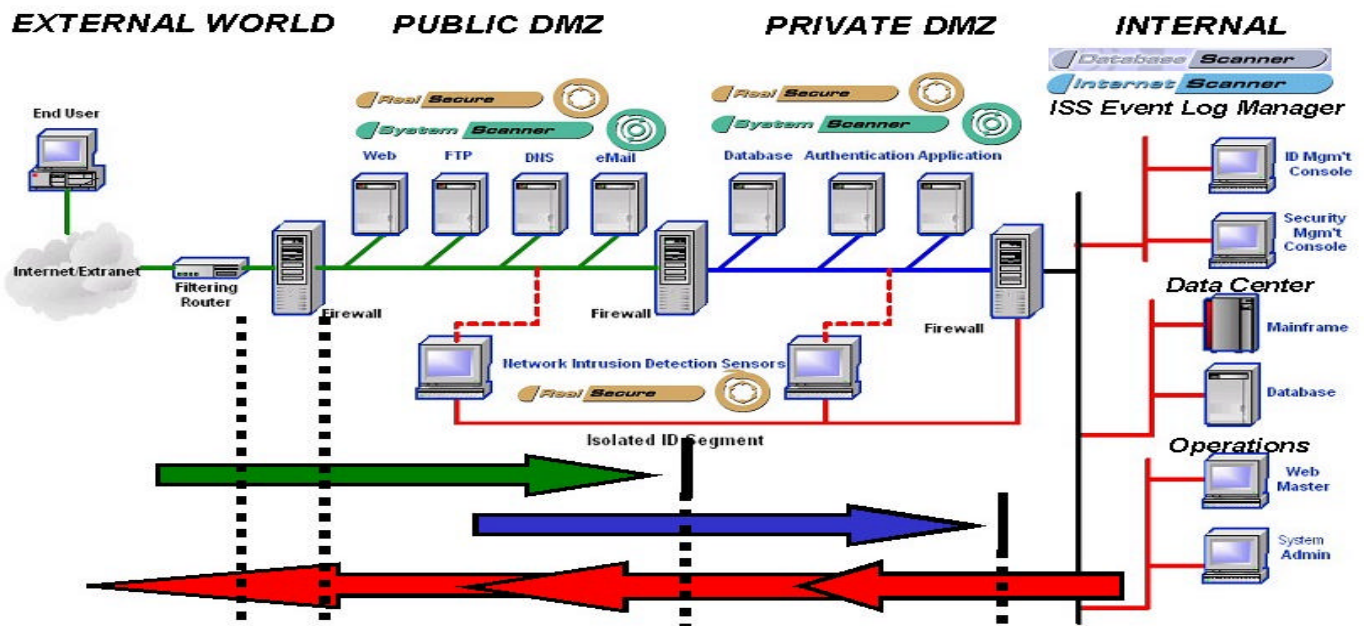
As organizations begin to exploit the benefit of the Internet and Web technologies, they are quickly realizing that there are inherent risks involved with connecting their networks to the Internet. This process of exposing valuable corporate systems and data significantly increases their risk of attack. As this transformation occurs, an organization's dependence on security, availability, and manageability significantly increases. Hence, security not only plays the role of protector but also of e-business enabler.

This paper describes a best-practices approach to information security for e-business. This approach includes an Internet/extranet architecture that uses network segmentation via multiple DMZ networks and an isolated IDS network segment, firewalls, authentication, encryption, intrusion detection, and vulnerability management. This approach enforces the rule of least privilege and manages risk, thereby preserving a reasonable level of data confidentiality, integrity, and availability. As risk cannot be eliminated, this paper also defines a comprehensive approach to incident management and response that enables organizations to effectively address this residual risk.

Organizations have traditionally spent significant amounts of time, effort, and money to protect their physical assets. As these assets move from the physical to the electronic, similar amounts of time, effort, and money must be applied to information security.

Addendum A

The following logical diagram represents an example of deployment of ISS Technology that will enable organizations to manage both the network and host security of the architecture. It is important to note that the ISS products such as RealSecure™ and System Scanner™ that require distributed installations support the communication controls and Network Address Translation (NAT) performed by the firewall(s) as detailed herein (the management consoles initiate all connections) and use very strong encryption and public/private key authentication mechanism that mitigates the risk of brute-force attacks. Although this diagram identifies three firewalls, these firewalls are intended to provide more of a pictorial representation of control rather than a requirement to have this number of firewalls. In a more physical network topological diagram, each of the segments (e.g., Public DMZ(s), Private DMZs, and Internal connectivity) can be constructed as separate interfaces on a single firewall (similar to the diagram above). The number of firewalls and redundancy (not depicted here) will be determined by performing an assessment of the business requirements.



RealSecure is a network and host based intrusion detection and response. RealSecure Sensors are composed of network sensors that are installed on the network on critical sub-networks and host sensors that are installed on critical production servers. A RealSecure Management Console centrally manages all sensors. Communication between the console and sensors is originated at the console and is strongly encrypted and authenticated using a public/private key authentication mechanism.

System Scanner is a host-based policy compliance and vulnerability management solution that also operates in a large distributed environment. Agents are installed on all production servers and controlled by a central management console. Communication is originated by the console and strongly encrypted and authenticated.

Internet Scanner™ is a network-based vulnerability assessment tool that is installed on a single system and can probe networks for vulnerabilities for a variety of network devices.

Database Scanner™ is a database security assessment solution. The Database Scanner assessment tool is installed on a single system and can assess Sybase, Oracle, and Microsoft SQL Server database systems distributed throughout the enterprise

ISS Event Log Manager™ is a log manager solution that enables organizations to systemically offload system logs from distributed systems to a central log management console for secure storage.

Authors' Biographies

MARC S. SOKOL

Marc is the National Manager of Business Development for the Internet Security Systems, Inc. (ISS) Emergency Response Service and joined ISS in December 1998. He is responsible for managing customer and partner relationships, managing ERS teams, conducting information security workshops, and assisting customers with the development of incident preparedness and response processes. Marc has over ten years experience in information systems and security, largely with the financial community. Marc has authored several white papers, including The Secure Deployment Methodology for Network-Based Intrusion Detection that has been accepted by major financial institutions and UK Government agencies as the best-practices approach to network intrusion technology deployment, A Methodology for Information Security Policy Development, and Baseline Electronic Warfare Modeling, a Guide to Performing Penetration Studies. Marc was also an integral developer of the ISS life-cycle security solution methodology and has appeared on BBC World News on the topic of Cyber defense. Prior to joining ISS, Marc was the Vice President of Information Security Global Distributed Infrastructure at Citigroup, with key responsibilities including managing perimeter security and firewall engineering, threat management and intrusion detection, vulnerability management and remediation, information security regulatory compliance, information security policy compliance and development, and incident response. Marc holds a cum laude Bachelors degree in Pre-law/Political Science from Brooklyn College.

DAVID A. CURRY

David joined ISS in January of 2000 as the Eastern Regional Manager of the ISS Emergency Response Service. David is responsible for managing technical incident management teams and participating in the development of the ISS-ERS business. David has authored three books: UNIX Systems Programming for SVR4 (O'Reilly & Associates), UNIX System Security: A Guide for Users and System Administrators (Addison-Wesley), and Using C on the UNIX System (O'Reilly and Associates). He is also the author of the well-known document "Improving the Security of Your UNIX System," and several popular public-domain software tools, all of which are widely distributed on the Internet.

Prior to joining ISS, David was one of the founders of IBM's Internet Emergency Response Service, where he served for four years as one of the lead members of the technical staff, and as business development manager. Before joining IBM, David worked in the UNIX systems programming and information security fields at Purdue University, SRI International, and the Research Institute for Advanced Computer Science at NASA Ames Research Center. David holds a Bachelor's degree in Computer Science from Purdue University.

Special thanks to Alan Fedeli, Alex Crepas, and David Nesom for their contributions and assistance with this paper.

About Internet Security Systems (ISS)

Internet Security Systems (ISS) is the leading global provider of security management solutions for the Internet. By providing industry-leading SAFEsuite® security software, ePatrol™ remote Managed Security Services, and strategic consulting and education offerings, ISS is a trusted security provider to its customers, protecting digital assets and ensuring safe and uninterrupted e- business. ISS' security management solutions protect more than 5,500 customers worldwide including 21 of the 25 largest U.S. commercial banks, 10 of the largest telecommunications companies and over 35 government agencies. Founded in 1994, ISS is headquartered in Atlanta, GA, with additional offices throughout North America and international operations in Asia, Australia, Europe, Latin America and the Middle East. For more information, visit the Internet Security Systems web site at <http://www.iss.net> or call 888-901-7477.

Internet Security Systems, ePatrol, Internet Scanner, System Scanner, Database Scanner and RealSecure are trademarks, and SAFEsuite a registered trademark, of Internet Security Systems, Inc. and ISS Group, Inc. All other companies and products mentioned are trademarks and property of their respective owners.