A Report for

National Institutes of Health

Security Architecture

18 July 2003

Engagement: 220441441

Table of Contents

Table of Contents

(Continued)

List of Figures

List of Tables

1.0 Introduction

Information security is critical to maintaining the confidentiality, integrity, and availability of NIH information and information systems. The goal of this security architecture is to provide a level of information security that is commensurate with risk in the most efficient manner possible.

1.1 Security Domain Team

The security architecture is the compilation of the work conducted by the Security Domain Team over a 12-week period. The Security Domain Team consisted of 10 subject matter experts from various Institutes and Centers (ICs). The individual contributors included:

- Vicky Ames, ORS
- Kevin Brownstein, CIT
- Kay Cornwell, NIGMS
- John Dvorak, CIT
- Dawn Farr, CIT
- Tim Ghebeles, NIAID
- Shawn Grimes, NIA
- Mike Iverson, OD/OIT
- Graham Logsdon, NLM
- Dan Sands, NCI

1.2 Scope

The scope of the first phase of the security architecture focuses on technical controls. These technical controls are divided into the following control areas:

- Identification and Authentication
- Access Control
- Confidentiality
- Integrity
- Monitoring and Analysis
- Availability.

1.3 NIH Enterprise Architecture Matrix

The NIH Enterprise Architecture Matrix is based on both the Federal Enterprise Architecture and the Zachman Framework. The NIH Enterprise Architecture Matrix allows the architecture process to be examined from numerous separate perspectives, See Table 1.

Table 1. NIH EA Matrix

	Data	Applications	Technology
Planner View	N/A	N/A	□ **Security Policy and Principles (Section 1.4).** High level rules for the application of security controls based on business drivers and risk analysis.
Owner View	N/A	N/A	□ **Security Model (Section 1.6).** High level rules for the association of classifications to security services.
Designer View	N/A	N/A	□ **Patterns (Section 2.0).** Definition of services/mechanisms within the context of a relationship between two logically separate domains.
Builder View	N/A	N/A	□ **Bricks (Section 3.0).** Definition of adoption timelines for various technologies, standards, products, and vendors.

► Relevant policies can be found located at the following address: http://irm.cit.nih.gov/security/sec_policy.html

1.4 Principles

The principles defined below are high level statements of the fundamental values that guide the information security architecture. Each principle should be universally accepted and should be stable so as to withstand changes in information security technologies, processes, and products. They should maintain a clear relevancy with policy changes in NIH programs and management approaches as well as reflect the general policy directions and framework of the Federal Government.

The principles are accompanied by rationales that explain their importance and business implications. While the statement of each principle should remain constant, the rationales and implications will evolve over time, as they respond to factors such as the current information management environment within NIH, internal initiatives, external forces, and changes in the NIH mission, vision, and strategic plan.

Table 2. Principles Developed by the Security Domain Team

Principle	Rationale
Principle #1: Information systems (including applications, computing platforms, data, and networks) should maintain a level of security that is commensurate with the risk and magnitude of the harm that could result from the loss, misuse, disclosure, or modification of information.	As a practical matter, perfect security cannot be achieved in any information system. Therefore, security controls should be applied to reduce risk to an acceptable level.
Principle #2: The selection of controls should be based on a risk analysis and risk management decision.	The selection of new controls should consider both the degree of risk mitigation provided by the control and the total cost to acquire, implement, and maintain the control.
Principle #3: Information system security should be built into systems from their inception rather than "bolted on" after system implementation.	The cost and complexity of adding security controls to a system after it is already in production is significantly greater.
Principle #4: Functional security requirements (that define the "what" a system or product does) should have associated assurance requirements (to define "how well it does it").	Security controls should be able to be reviewed or audited through some qualitative or quantitative means in order to ensure that risk is being maintained at acceptable levels.
Principle #5: Safeguards should be modular so that they may be removed or changed as the system and enterprise risk profile changes.	It is prudent to minimize the interdependence of controls so that controls can be easily interchanged or modified.
Principle #6: Selection of controls should consider the ability of the control to be applied uniformly across the NIH Enterprise and to minimize exceptions.	Achieving a standards-based environment should reduce operational costs, improve interoperability, and supportability.
Principle #7: The architecture should embrace the concepts of security domains and layered defense.	Compartmentalization localizes vulnerabilities and defense-in-depth establishes a series of imperfect countermeasures.
Principle #8: Controls should minimize the need for manual operation.	Manual operation can create vulnerabilities and cause disruption of service due to misinterpretation and misconfiguration.
Principle #9: Controls should not impose unreasonable constraints or operate in a manner that causes unreasonable response.	Controls should provide only the required level of protection, alerting, and response.
Principle #10: Controls should have the capability to be shutdown gracefully and restored automatically to the conditions prior to shutdown.	Controls should be highly available to minimize periods of vulnerability.
Principle #11: Controls will default to a lack of permission.	This approach reduces the number of points of vulnerability.

1.5 Drivers

Generally, a business driver is an external pressure or initiative that impacts the way an organization responds to the market. These are either direct pressures (for example, entrance into a new market) or indirect pressures (for example, new competition in an existing market). Technical drivers are those factors that are changing the technology environment for not only the NIH, but for industry in general.

This section describes the drivers that are impacting the NIH security architecture.

Figure 1. NIH Technical and Business Drivers

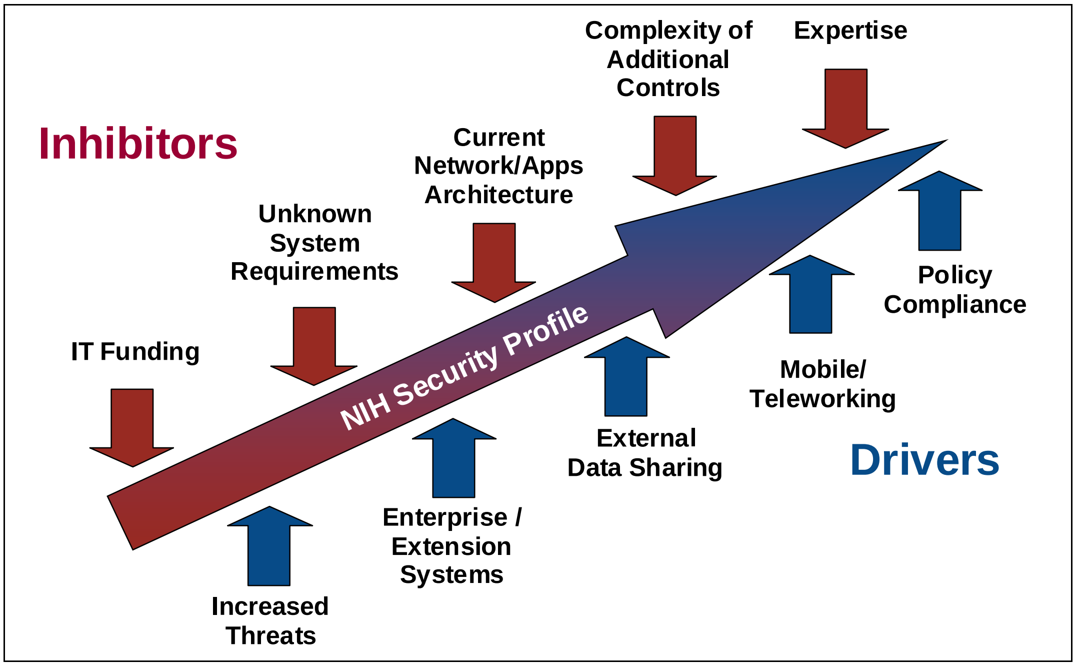

Source: Gartner, 2003

Figure 1 illustrates that there are a number of business and technical drivers that are accelerating the increase in the overall NIH security posture. These drivers include:

- **An increase in the number of threats.** This increase is due to both increases in the number of general threat agents such as viruses and the number of threats to high profile government organizations by hactivists or foreign governments.
- **An increase in enterprise systems and extension systems.** This increases the risk to NIH because systems that store, process, or transmit sensitive information are being made available to a larger number of internal and external users.
- **External Data Sharing.** The need to share sensitive information inside and outside of NIH increases the vulnerability to unintended disclosure.
- **Mobile Computing/Teleworking.** The traditional enterprise network is being extended to mobile and teleworking employees. As the network is extended to remote uses the ability to enforce standard security controls becomes increasing difficult.
- **Policy Compliance.** Government organizations are now required through federal regulations to comply with Department and Federal information security policy.

Just as there are factors driving an increase in NIH's security posture there are several factors, that must be effectively managed, that inhibit NIH's ability to increase its posture. These inhibitors include:

- **IT Funding.** The overall NIH IT budget is expected to decrease over the next several years. Since NIH has not created a uniformly secure environment across the IC it should be expected that initial investments would need to be made.
- **Unknown System Requirements.** Most systems within NIH have not been classified and a certification and accreditation process has not yet been fully implemented. This means that the security requirements of many systems are undocumented or unknown.
- **Current network and applications architecture.** Currently, the network and applications architecture has not been fully consolidated. The implication is that security components may need to be deployed in a distributed manner.
- **Complexity of additional controls.** As NIH adds more security controls to keep pace with increasing threats and vulnerabilities, the complexity of the overall environment will increase. This also due to the fact that some technologies that are not fully mature may need to be implemented.
- **Expertise.** Many cutting edge security technologies have a steep learning curve and developing and retaining expertise can be difficult.

1.6 Security Model

The Department of Health and Human Services Information Security Handbook defines the information security classification scheme for the Department including NIH. This classification scheme is critical to achieving Principle #1 *"Information systems (including applications, computing platforms, data, and networks) will maintain a level of security that is commensurate with the risk …*" Therefore, in order to apply the proper security controls, a system's level of sensitivity and criticality must be understood. Table 3 presents the levels of classification and a description of each level.

Table 3. Department of Health and Human Services Security Classifications

Sensitivity (Confidentiality)	Criticality (Integrity and Availability)
Level 1: Low Sensitivity. Data that require minimal protection. Threats to these data are minimal, and only minimal precautions to protect the data need to be taken at the user site.	**Level 1: Low Criticality.** AISs with data processing capabilities that require minimal protection. These include AISs that, in the event of alteration or failure, would affect the organization minimally or could be replaced with a minimum of staff time or expense.
Level 2: Moderate Sensitivity. Data that have some importance to the Agency and which must be protected against such acts as malicious destruction. Data are most often collected for analytical purposes.	**Level 2: Moderate Criticality.** AISs with data processing capabilities that are considered important but not critical to the internal management of an organization and/or the Department.
Level 3: High Sensitivity. The most sensitive unclassified data (other than national security interests). Requires the greatest number and most stringent security safeguards at the user level.	**Level 3: High Criticality.** AISs with data processing capabilities that are considered critical to the organizations they support and/or the Department.
Level 4: High Sensitivity and National Security Interest. National security classified information and all databases that contain other sensitive, but unclassified information, the loss of which could adversely affect national security interests.	**Level 4: High Criticality and National Security Interest.** AISs with data processing capabilities that are considered critical to the well being of the nation.

► For more information on DHHS Security Level Designations go to:
http://irm.cit.nih.gov/policy/Introduction%20to%20Security%20Levels

Based on the DHHS security level designations a security model has been developed that defines, at a high level, the various security services that are required at the client, network, and server. The various mechanisms (that is, standards, technologies, and products) that provide the required security services are not defined by this model, but are defined by the patterns and bricks presented later in this document.

In fact, the application of specific mechanisms will be dependent on a risk analysis that considers both degrees of risk mitigation and cost. For example, in some cases a business partner boundary (see Pattern 3: Business Partner Boundary) will implement a firewall to provide logical access control, in other cases a router with access control lists may be used. However, in both cases the logical access control defined by the security model and the boundary pattern is provided. Figure 2 is a graphical depiction of the security model.

Figure 2. Information Security Model

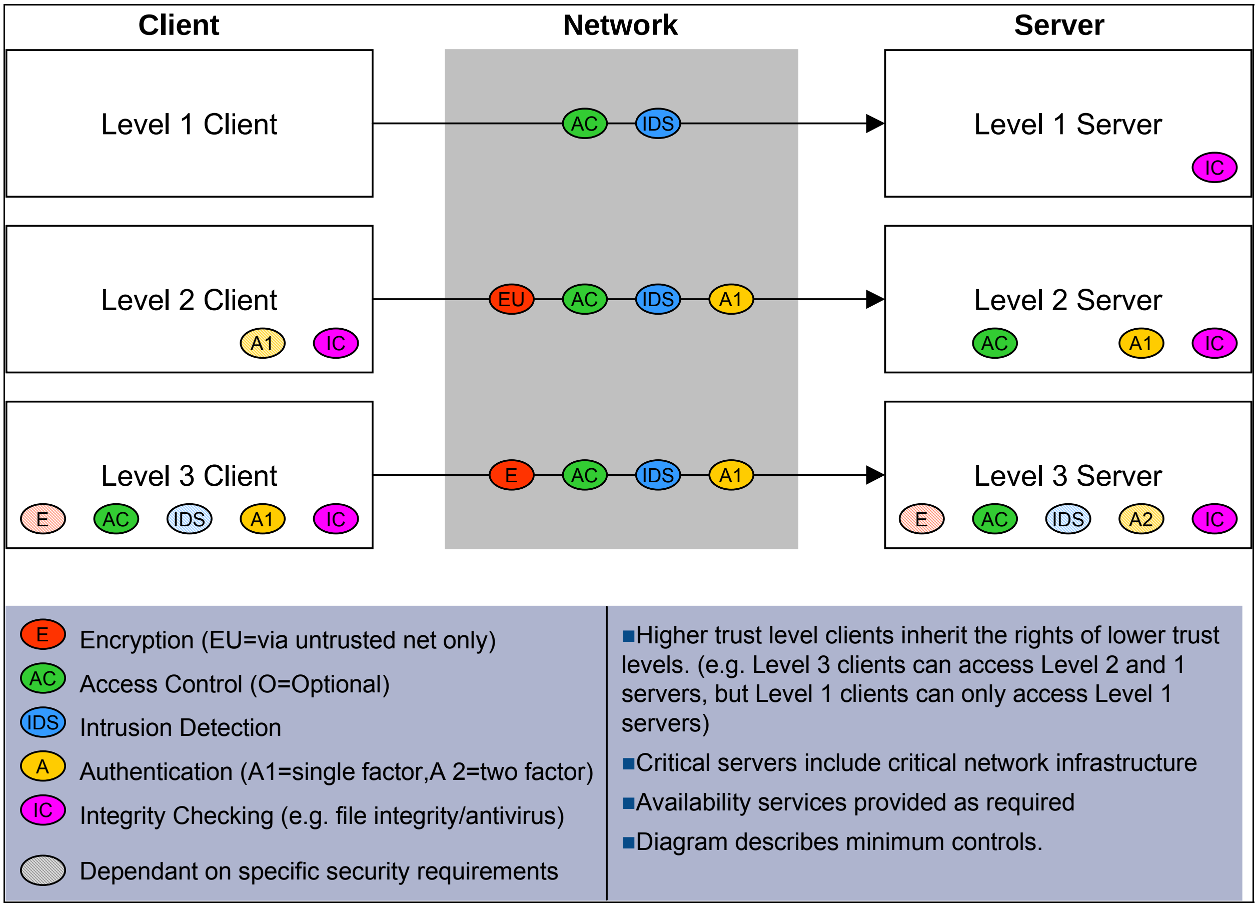

- E Encryption (EU=via untrusted net only)
- AC Access Control (O=Optional)
- IDS Intrusion Detection
- A Authentication (A1=single factor,A 2=two factor)
- IC Integrity Checking (e.g. file integrity/antivirus)
- ○ Dependant on specific security requirements

- Higher trust level clients inherit the rights of lower trust levels. (e.g. Level 3 clients can access Level 2 and 1 servers, but Level 1 clients can only access Level 1 servers)
- Critical servers include critical network infrastructure
- Availability services provided as required
- Diagram describes minimum controls.

Source: Gartner, 2003

2.0 Patterns

Figure 3 presents the security architecture patterns in an overall context. Each pattern is a definition of the security services and mechanisms that must be in place at a boundary. A boundary represents a point of logical separation between two domains. Each domain may represent a network, a server, a client, or a collection of networks, servers, and clients.

Figure 3. NIH Information Security Boundary Patterns

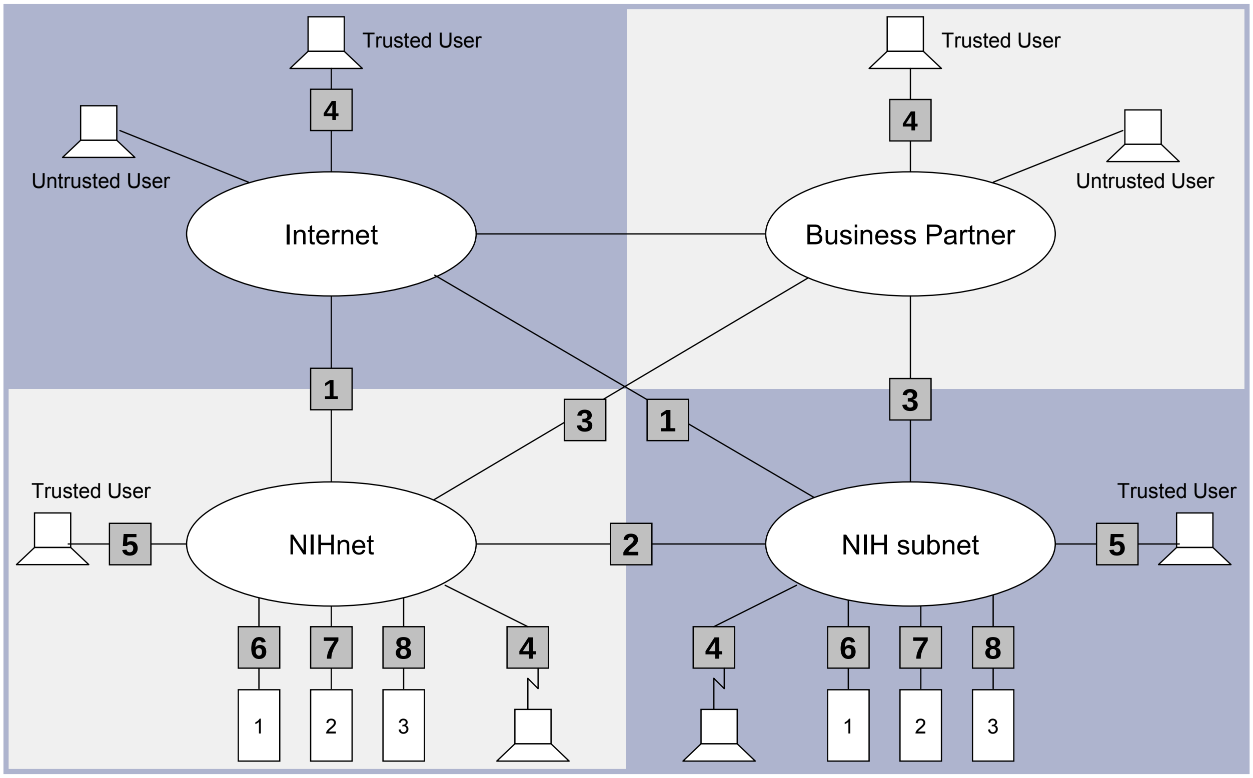

Source: Gartner, 2003

- Pattern 1: Internet Boundary
- Pattern 2: NIHnet/ICnet Boundary
- Pattern 3: Business Partner Boundary
- Pattern 4: Remote Access Boundary
- Pattern 5: Trusted User Boundary
- Pattern 6: Level 1 System Boundary
- Pattern 7: Level 2 System Boundary
- Pattern 8: Level 3 System Boundary

The following sections describe each pattern and define the specific security services required for each boundary. Each pattern includes a description of the benefits and limitations of the pattern. In a security context, the benefits can be regarded as risk mitigation and limitations can be regarded as residual risk. Descriptions of each of the security services can be found in Appendix A – Glossary of Terms.

The requirement for each security service is specified as mandatory, dependent or optional. Definitions are as follows:

- **Mandatory**—The specified security service must be in place in every instance of the boundary.
- **Dependent**—The specified security service may or may not be required depending on the sensitivity or criticality of the information that traverses the boundary. The security model should be applied in these cases to determine if a particular security service is required at a boundary.
- **Optional**—The specified security service is not explicitly required by the security model but may be employed at the discretion of a system/domain owner.

2.1 Pattern 1: Internet Boundary

2.1.1 Context

This pattern defines the boundary architecture between NIHnet and the public Internet. This boundary is where the majority of external access to internal systems crosses. In addition, all internal access to external systems crosses this boundary. This boundary is where the majority of external threats will need to be detected and prevented. Unique to Boundary Pattern 1 is the use of content filtering on inbound traffic to support enforcement of NIH Internet "acceptable use" policies.

2.1.2 Solution

Table 4. Internet Boundary Services (July 2003)

Control	Service	Requirement
Access Control	Content-Dependant Access Control	Mandatory
	Rules-based (ACL) Access Control	Mandatory
User Authentication	Single Factor Authentication	Dependent
Audit/Analysis	Intrusion Detection	Mandatory
Integrity	Malicious Code Management	Mandatory
Confidentiality	Encryption	Dependent
Availability	High Availability	Mandatory

2.1.3 Benefits

This boundary provides a consolidated ingress and egress point where a variety of security mechanisms can be implemented efficiently. This boundary is where the majority of external attacks occur.

2.1.4 Limitations

This boundary provides no protection from internal threats to internal systems.
The security mechanisms in this boundary must also be in place for alternate network routing paths.

2.2 Pattern 2: NIHnet/ICnet Boundary

2.2.1 Context

This boundary is between the NIHnet backbone and any IC sub-network (ICnet). An ICnet may be a single logical sub-network or may be several sub-networks spread across NIH and connected via the NIHnet backbone.

2.2.2 Solution

Table 5. NIHnet/ICnet Boundary Services (July 2003)

Control	Service	Requirement
Access Control	Content-Dependant Access Control	Optional
	Rules-based (ACL) Access Control	Mandatory
User Authentication	Single Factor Authentication	Mandatory
Audit/Analysis	Intrusion Detection	Mandatory
Confidentiality	Encryption	Dependent
Availability	High Availability	Mandatory

2.2.3 Benefits

This boundary provides containment of incidents, protection from internal threats, and/or provides the additional controls necessary to protect sensitive or critical IC resources at a consolidated point.

2.2.4 Limitations

Currently, only a few ICs have connectivity to NIHnet in a manner that allows consolidated implementation of security mechanisms. Network architecture changes or larger investments may be required by these ICs.

2.3 Pattern 3: Business Partner Boundary

2.3.1 Context

This boundary is between NIHnet and a business partner. The business partner domain is assumed to be untrusted primarily because it is not under direct NIH management and NIH cannot control changes in the business partner network that may increase risk to NIH. However, the business partner domain may be host to trusted and untrusted users that will need to access systems on NIHnet. Similarly, users on NIHnet will need to access systems within the business partner domain.

2.3.2 Solution

Table 6. Business Partner Boundary Services (July 2003)

Control	Service	Requirement
Access Control	Content-Dependant Access Control	Optional
	Rules-based (ACL) Access Control	Mandatory
User Authentication	Single Factor Authentication	Mandatory
Audit/Analysis	Intrusion Detection	Mandatory
Integrity	Malicious Code Management	Mandatory
Confidentiality	Encryption	Dependent
Availability	High Availability	Dependent

2.3.3 Benefits

Treating business partner domains as untrusted sites protects NIHnet even if the business partner changes their environment and accepts new risks.

2.3.4 Limitations

There are situations where memorandums of agreement can be established to mitigate the risks associated with lack of change control in the business partner domain(s). In these cases it may be possible to treat business partner connections the same as a Pattern 2: NIHnet/ICnet boundary. However, most of the same controls will still be required to provide monitoring and containment of attacks.

2.4 Pattern 4: Remote Access/Wireless Boundary

2.4.1 Context

The remote access boundary applies to all forms of remote access including Internet or business partner VPN, dial-in remote access, and wireless. By definition, the remote access boundary pattern assumes that an untrusted network (i.e. a network that is not owned or managed by NIH or is not via NIH managed VPN and encryption) is being traversed for by a trusted client (i.e. a client that implements NIH managed or specified security services) to trusted server communications. Even wireless local area networks are considered remote access since the medium that is traversed between the client and the wireless network access point is considered untrusted, as it is inherently vulnerable to remote monitoring and traffic that is designed to penetrate or attack NIH resources.

2.4.2 Solution

Table 7. Remote Access/Wireless Boundary Services (July 2003)

Control	Service	Requirement
Access Control	Content-Dependant Access Control	Optional
	Rules-based (ACL) Access Control	Mandatory
User Authentication	Single Factor Authentication	Mandatory (Network) / Mandatory (Client)
Audit/Analysis	Intrusion Detection	Mandatory
Integrity	Malicious Code Management	Mandatory (Client)
Confidentiality	Encryption	Mandatory
Availability	High Availability	Optional

2.4.3 Benefits

The benefits of a single model for evaluating and protecting all remote access is that as new methods of remote access emerge or as use of existing remote access changes at NIH there will be a consistent approach to providing security for these environments.

2.4.4 Limitations

This boundary pattern assumes that the users and devices communicating across this boundary are trusted. However, there are scenarios where untrusted users may need access to NIHnet. For example: a presentation given by a visitor who needs Internet access for the purposes of demonstrating a web site. In these cases, visiting users will be restricted to those services available to external (that is, Internet resident) untrusted users. That is use of the Internet (subject to NIH content-based restrictions) and access to NIH Level 1 servers.

2.5 Pattern 5: Trusted User Boundary

2.5.1 Description

This boundary pattern addresses the controls required for a trusted client to locally (that is, physical connection to an NIH managed network within NIH managed facilities) access the NIH network and server resources. Some of the controls will be resident exclusively on the client, some of the controls will be both a client and network service (for example, authentication).

2.5.2 Solution

Table 8. Trusted User Boundary Services (July 2003)

Control	Service	Requirement
Access Control	Content-Dependant Access Control	Optional
	Rules-based (ACL) Access Control	Mandatory
User Authentication	Single Factor Authentication	Mandatory (Network) / Mandatory (Client)
Audit/Analysis	Intrusion Detection	Dependant
Integrity	Malicious Code Management	Mandatory
Confidentiality	Encryption	Dependant
Availability	High Availability	Optional

2.5.3 Benefits

Controls defined at the client level and based on system classification are important because the servers and data that these clients access can be intentionally or unintentionally stored on the client, turning the client into a potential server of that sensitive information.

2.5.4 Limitations

Some systems that connect to NIH networks are systems that are not under the direct management of NIH. Therefore, in some cases a Memorandum of Agreement with a user or another organization will need to be developed to ensure that proper client security controls are implemented.

Users who do not satisfy the mandatory service requirements below are considered untrusted users are an therefore have restricted access to Level 1 systems only.

2.6 Pattern 6: Level 1 System Boundary

2.6.1 Context

This boundary applies to Level 1 servers. These servers are generally used to provide information to external organizations and to the general public. The data resident on these systems is not restricted in its dissemination but it is still important that the information is accurate and available.

2.6.2 Solution

Table 9. Level 1 System Boundary Services (July 2003)

Control	Service	Requirement
Access Control	Content-Dependant Access Control	Optional
	Discretionary Access Control	Optional
	Rules-based (ACL) Access Control	Mandatory (Network)
User Authentication	Single Factor Authentication	Optional
Audit/Analysis	Intrusion Detection	Mandatory (Network) / Optional (Server)
Integrity	Data Integrity/Malicious Code	Mandatory
Confidentiality	Encryption	Optional

2.6.3 Benefits

The benefit of this pattern is that only those mechanisms that are required to protect the integrity and availability of the data are mandatory, thereby minimizing the cost and complexity to protect these systems.

2.6.4 Limitations

This boundary does not address data aggregation. That is to say the notion that multiple sources of public information can be aggregated to create or make inferences about sensitive information.

2.7 Pattern 7: Level 2 System Boundary

2.7.1 Context

This boundary addresses Level 2 systems. These systems are generally available to NIH employees and business partners who are involved in day to day NIH business processes. These systems may physically reside within NIH or in some cases reside within business partner domains.

2.7.2 Solution

Table 10. Level 2 System Boundary Services (July 2003)

Control	Service	Requirement
Access Control	Role-based Access Control	Dependent
	Discretionary Access Control	Mandatory
	Rules-based Access Control	Mandatory (Network)
User Authentication	Single Factor Authentication	Mandatory
Audit/Analysis	Intrusion Detection	Mandatory(Network) / Optional (Server)
Integrity	Malicious Code/Data Integrity	Mandatory (Client) Optional (Server)
Confidentiality	Encryption	Dependent

2.7.3 Benefits

The benefit of this pattern is a higher level of security is required for the system even though the system is intended for internal or business partner use only. These controls minimize the risk associated with intentional or unintentional disclosure or disruption of service by insiders.

2.7.4 Limitations

This pattern cannot prevent sensitive information that has been retrieved from the server from being disclosed or altered.

2.8 Pattern 8: Level 3 System Boundary

2.8.1 Context

This pattern addresses the requirements of the most sensitive systems within NIH. These systems contain information that is subject to HIPAA and privacy act regulations. These systems contain also business information of a critical nature such as a financial management data where integrity and availability requirements are high.

2.8.2 Solution

Table 11. Level 3 System Boundary Services (July 2003)

Control	Service	Requirement
Access Control	Role-based Access Control	Dependent
	Rules-based Access Control	Mandatory (Network)
	Discretionary Access Control	Mandatory
User Authentication	Single Factor Authentication	Mandatory
Audit/Analysis	Intrusion Detection	Mandatory(Network) / Dependent (Server)
Integrity	Malicious Code/Data Integrity	Mandatory
Confidentiality	Encryption	Dependant (Network) / Dependent (Server)

2.8.3 Benefits

The number of these types of systems is relatively low so the cost of providing the additional security control can be minimized.

2.8.4 Limitations

As with Pattern 7, this pattern cannot prevent sensitive information that has been legitimately retrieved from the server from being disclosed or altered.

3.0 Bricks

3.1 Brick 1: Identification and Authentication

Authenticated identities are the basis for many other information security services. Therefore, NIH needs to:

- Verify user identity as the basis for access control to NIH resource
- Control individual user access to the resources and services provided by those systems
- Create an audit trail of individual user access or attempted access to those systems, resources and services.

Authentication services are crucial to access control and auditing services. If users' identities are not properly authenticated, NIH has no assurance that access to resources and services are properly controlled. No matter how well managed NIH's access control services; everything hinges on the true identity of the user. In most situations, User ID and password combinations will provide an appropriate level of security if the User ID and password conform to NIH policy. However, NIH will implement stronger authentication for Enterprise users with high system privileges—that is, system, network and security administrators.

Brick 1. Identification and Authentication (July 2003)

Baseline Environment (Today)		**Tactical Deployment (0-2 Years)**	**Strategic (2-5 Years)**
■ User ID/ Password ■ MAC Address (LAN and WLAN) ■ Repository ❑ Active Directory ❑ Novell NDS ❑ RADIUS Database ❑ Cisco ACS ■ PKI ■ Tokens ❑ RSA SecureID ■ Smartcards	■ Protocols ❑ Kerberos ❑ TACACS + ❑ RADIUS ❑ Cisco LEAP ❑ LDAP ❑ SSH ❑ SSL ❑ IKE ■ Web Single Sign On ❑ Netegrity Sightminder	■ User ID/Password ■ Tokens ■ Web Single Sign On (SSO) ■ Kerberos ■ RADIUS ■ SSL ■ LDAP	■ TBD
Retirement (Technology to eliminate)		**Containment (No new deployments)**	**Emerging (Technology to track)**
■ TBD		■ TACACS+ ■ WEP ■ LEAP	■ PKI ■ Multiplatform SSO ■ Biometrics ■ Smartcards ■ EAP-TTLS ■ PEAP

Comments	
Relevant Standards ■ IETF RFC2058 Remote Authentication Dial-In User Service (RADIUS) ■ RFC 1510 Kerberos Authentication Service ■ FIPS PUB 196 Entity Authentication Using Public Key Cryptography ■ IEEE 802.10: Interoperable LAN/MAN Security (SILS)	**Relevant Policies** ■ CIT Security Handbook (NIH) ■ Guidance for Selecting Good Passwords (NIH) ■ Policy on Passwords (NIH) ■ Active Directory (DHHS) ■ DHHS AISSP Handbook (DHHS) ■ Usage of Persistent Cookies (DHHS) ■ Public Key Infrastructure Certification Authority (DHHS)

3.2 Brick 2: Access Control

Logical access control within NIH is provided at the network, operating system, and application level.

- **Network Access Control.** Network access controls can be provided by a variety of mechanisms both alone and in combination. However, the primary method of providing network access control in an enterprise environment is via a firewall. By 1Q04, Gartner predicts that more than 50 percent of Fortune 1000 enterprises will have distributed firewalls internally.
- **System Access Control.** Access control can also be provided by the client or server operating system. Host access control can also be provided at the operating system level via third party products that are designed to enhance an operating system's native access control facilities.
- **Application Access Control.** Application access control can be provided by either the underlying Data Base Management System (DBMS) or by the application itself.
- **Content Filtering.** Access control can also be based on content or sites. The motivation to block certain content or sites is driven by NIH acceptable use policy.

Brick 2. Access Control (July 2003)

Baseline Environment (Today)		Tactical Deployment (0-2 Years)	Strategic (2-5 Years)
■ Firewalls ❑ Cisco PIX ❑ Gauntlet ❑ Lucent ❑ BorderManager ❑ Enterasys ❑ Checkpoint ❑ Netscreen ■ Other Network Access Control ❑ MAC Address ACLs ❑ Network Address Translation ❑ VLAN ❑ Router Access Control Lists ❑ SSID ❑ Domain Blocking ❑ VPN ❑ IP Tables	■ Repository ❑ Active Directory ■ System Access Control ❑ SAMBA ❑ IBM Host Access Class Library ❑ Pelican ❑ TCP/IP Wrappers ❑ Sudo ❑ Okena Stormwatch ❑ Citrix CSG ■ Application Access Control ❑ Role-based access control ❑ DBMS ■ Content Filtering ❑ Websense	■ TBD	■ TBD
Retirement (Technology to eliminate)		**Containment (No new deployments)**	**Emerging (Technology to track)**
■ Pelican		■ Gauntlet (Application proxy requirements only) ■ Lucent ■ IP Chains	■ Host Based Firewall ■ Intrusion Prevention

Comments	
Relevant Standards ■ IEEE 802.10: Interoperable LAN/MAN Security (SILS) ■ ISO/IEC 10164-9 System Management: Objects and Attributes for Access Control ■ NIST Generally Accepted Principles and Practices for Secure Information Technology Systems	**Relevant Policies** ■ CIT Security Handbook (NIH) ■ NIH FTP Server Policy ■ NIH Guidance on File Sharing Security ■ NIH NetBIOS Security Policy (NIH) ■ NIH NIHnet/Firewall Policy (NIH) ■ NIH SMTP Server Policy (NIH) ■ NIH SQL Server Policy (NIH) ■ NIH Web Server Policy (NIH) ■ DHHS AISSP Handbook (DHHS)

3.3 Brick 3: Confidentiality

The primary method of protecting confidentiality of information is via encryption. In addition to sensitive business data, there is also data about the network and systems themselves that need to be encrypted in order to prevent attacks.

Brick 3. Confidentiality (July 2003)

Baseline Environment (Today)		**Tactical Deployment (0-2 Years)**	**Strategic (2-5 Years)**
■ Key Size □ 168 bit □ 128bit □ 40bit □ 56 bit ■ Algorithms □ 3DES □ MD5 ■ Database Encryption ■ Message encryption □ PGP □ Microsoft	■ Transport Encryption □ IPSEC □ WEP □ SSL □ SSH □ Blackberry Transport Encryption ■ LZS compression ■ Data Link Encryption ■ File encryption □ PGP	■ Baseline	■ TBD
Retirement (Technology to eliminate)		**Containment (No new deployments)**	**Emerging (Technology to track)**
■ 40 bit ■ LZS compression (retired as an encryption mechanism only)		■ WEP ■ 56 bit	■ AES ■ 802.11I ■ S-HTTP ■ S/MIME
Comments			
Relevant Standards ■ FIPS 180-2, August 2002, SECURE HASH STANDARD ■ FIPS 46-3, October 1999, Data Encryption Standard (DES) ■ FIPS 197 November 2001, Advanced Encryption Standard ■ IEEE 802.10: Interoperable LAN/MAN Security (SILS) ■ IETF ID Combined SSL/PCT Transport Layer Security Protocol ■ IETF ID Secure HyperText TP Protocol (S-HTTP) ■ IETF ID SMIME Cert Handling ■ IETF ID SMIME Message Specification ■ Security Architecture for the Internet Protocol (RFC 2401)		**Relevant Policies** ■ Guidance for Securing Data on Portable Systems (NIH) ■ Sanitization Policy (NIH) ■ Security Guidelines for NIH Remote Access Users (NIH) ■ IT Security for Remote Access (DHHS) ■ Public Key Infrastructure Certification Authority (DHHS)	

3.4 Brick 4: Integrity

- **Anti-Virus.** An effective anti-virus architecture uses a multi-tier (that is, desktop, server, and gateway) approach and is not necessarily reliant on a single vendor solution. The gateway tier can be implemented at the firewall, the SMTP gateway, the SMTP relay, or a combination of all three. Using a combination of techniques at the gateway level is prudent given the frequency and impact of malicious code attacks. NIH currently implements a multi-tier anti-virus architecture.
- **Configuration Management.** Configuration management is the basis for all other management capabilities and is a critical aspect of maintaining confidentiality, integrity, and availability. Change management and software control and distribution must be properly integrated with a comprehensive configuration management system.
- **File Integrity Checking.** File integrity checking is used to detect and correct unauthorized changes to a file or database.

Brick 4. Integrity (July 2003)

Baseline Environment (Today)		**Tactical Deployment (0-2 Years)**	**Strategic (2-5 Years)**
■ Network Integrity ❑ Anti-spoofing filters ■ Anti-Virus ❑ Norton Anti-virus ❑ Norton Command Center ❑ McAfee Anti-Virus ❑ McAfee ePolicy Orchestrator ❑ Sybari Antigen for Exchange ❑ Symantec Virus Scan	■ File Integrity Checking ❑ Samhain File Integrity ❑ Tripwire ■ Digital Signature ■ Configuration Management ❑ Ecora ❑ Peregrine IND ❑ Bindview ❑ HFNetChk Pro ❑ Update Expert ❑ Alteris	■ Baseline ■ Bluesocket Secure Gateway ■ Tripwire	■ TBD
Retirement (Technology to eliminate)		**Containment (No new deployments)**	**Emerging (Technology to track)**
■ TBD		■ Peregrine IND	■ TBD
Comments			
Relevant Standards ■ IEEE 802.10: Interoperable LAN/MAN Security (SILS) ■ ISO/IEC 10736 Information Technology - Telecommunications and Information Exchange Between Systems - Transport Layer Security Protocol (TLSP) ■ ISO/IEC 11577 Information Technology - Telecommunications and Information Exchange Between Systems - Network Layer Security Protocol (NLSP) ■ Security Architecture for the Internet Protocol (RFC 2401) ■ NIST SP 800-40 Procedures for Handling Security Patches, September 2002		**Relevant Policies** ■ CIT Security Handbook (NIH) ■ Guidance for Securing Data on Portable Systems (NIH) ■ Procedures for Handling Unwanted E-mails (NIH) ■ DHHS AISSP Handbook (DHHS) ■ Prevention, Detection, Removal and Reporting of Malicious Software (DHHS)	

3.5 Bricks 5,6,7: Monitoring and Analysis

- **Vulnerability Analysis.** Internet-based attack tools are becoming increasingly sophisticated and increasingly easy to use. NIH's network could contain vulnerabilities that attackers can exploit to gain access, even when NIH has secured the network perimeter with firewalls and intrusion detection systems. In order to proactively find and plug such holes NIH will require the use of both vulnerability assessment products and vulnerability assessment services.
- **System Monitoring and Logging.** Identifying and reacting to security incidents in real-time requires comprehensive system and network monitoring, Furthermore the ability to aggregate alarms and other information from disparate systems is necessary to correlate events and identify an incident.

Brick 5. Event Monitoring and Analysis(July 2003)

Baseline Environment (Today)		**Tactical Deployment (0-2 Years)**	**Strategic (2-5 Years)**
■ Event Monitoring ❑ Computer Associates TNG ❑ HP Openview ❑ HP ITO ❑ Fluke Optview ❑ Fluke Link Analyzer ❑ Fluke Network Inspector ❑ Quest Software Big Brother ❑ Open NMS ❑ NetIQ ❑ Ipswitch What's Up Gold ❑ Deepmetrix ipMonitor	■ Log Monitoring ❑ Microsoft Operations Manager ■ Log Analysis ❑ Envision ❑ Central Syslog Facility ❑ Router/switch logging ❑ Remote syslog ❑ OS Logging ❑ Microsoft Operations Manager	■ TBD	■ TBD
Retirement (Technology to eliminate)		**Containment (No new deployments)**	**Emerging (Technology to track)**
■ None		■ None	■ Event Correlation
Comments			
Relevant Standards ■ ISO/IEC 10164-4 System Management: Alarm Reporting Function ■ ISO/IEC 10164-5 System Management: Event Report Management Function ■ ISO/IEC 10164-7 System Management: Security Alarm Reporting Function ■ ISO/IEC 10164-8 System Management: Security Audit Trail Function		**Relevant Policies** ■ CIT Security Handbook (NIH) ■ Glossary of NIH Incident Handling Terms (NIH) ■ Limited Authorized Personal Use of NIH Information Technology (IT) Resources (NIH) ■ Incident Handling Guidelines (NIH) ■ NIH Incident Handling Procedures Text and Diagram (NIH) ■ Establishing an Incident Response Capability (DHHS) ■ OMB Circular A-130 (Federal)	

Brick 6. Vulnerability Analysis (July 2003)

Baseline Environment (Today)	**Tactical Deployment (0-2 Years)**	**Strategic (2-5 Years)**
■ Third Party Vulnerability Assessments ■ Vulnerability Analysis ❑ LANGuard ❑ Enterprise Security Manager ❑ Nessus scanner SARA Scan ❑ MS Baseline security Analyzer ❑ Wireless Sniffer/Site Surveys ❑ Harris Stat Analyzer ❑ ISS Internet Scanner ❑ eEye Retina ■ Port scanner ❑ NMAP ■ Vulnerability Remediation ❑ Citadel Hercules	■ TBD	■ TBD
Retirement (Technology to eliminate)	**Containment (No new deployments)**	**Emerging (Technology to track)**
■ None	■ None	■ TBD

Comments	
Relevant Standards ■ ISO/IEC 10181-7 Security Audit Framework ■ NIST SP 800-51 Use of the Common Vulnerabilities and Exposures (CVE) Vulnerability Naming Scheme, September 2002 ■ NIST SP 800-24 PBX Vulnerability Analysis, August 2000	**Relevant Policies** ■ OMB Circular A-130 (Federal)

Brick 7. Intrusion Detection (July 2003)

Baseline Environment (Today)	**Tactical Deployment (0-2 Years)**	**Strategic (2-5 Years)**
■ Network Intrusion Detection ❑ Snort IDS ❑ ISS Real Secure ❑ NAI Intruvert ■ Host Intrusion Detection ❑ Tripwire ❑ Black Ice Software ■ Event correlation ■ Honeypot ■ Intrusion Prevention ❑ Okena Stormwatch ❑ INtruvert	■ TBD	■ TBD
Retirement (Technology to eliminate)	**Containment (No new deployments)**	**Emerging (Technology to track)**
■ None	■ None	■ Checkpoint NG

Comments	
Relevant Standards ■ ISO/IEC 10164-4 System Management: Alarm Reporting Function ■ ISO/IEC 10164-5 System Management: Event Report Management Function ■ ISO/IEC 10164-7 System Management: Security Alarm Reporting Function ■ ISO/IEC 10164-8 System Management: Security Audit Trail Function ■ NIST SP 800-31 Intrusion Detection Systems (IDS), November 2001	**Relevant Policies** ■ CIT Security Handbook (NIH) ■ Glossary of NIH Incident Handling Terms (NIH) ■ Incident Handling Guidelines (NIH) ■ NIH Incident Handling Procedures Text and Diagram (NIH) ■ Establishing an Incident Response Capability (DHHS)

4.0 Next Actions

The following actions for the Security Domain Team have been identified as a follow on to the initial phase of enterprise architecture development project:

- Validate consolidation effort against the security architecture
- Perform transition planning—This will include a full analysis of targets for tactical deployment, retirement, containment, strategic, and emerging
- Evaluate impact of new HHS information security policies on the security architecture
- Develop patterns and bricks for high availability to include areas such as load balancing, fault tolerance, redundancy, replication, backup and recovery, clustering, and storage area networks.

Appendix A—Glossary of Terms

Appendix A—Glossary of Terms

Term	Definition
Anti-Virus Software	Anti-Virus software is a class of program that searches your hard drive and floppy disks for any known or potential viruses.
Application Encryption	Method of security by encrypting the application layer of the transmission
Biometrics	Biometrics is the science and technology of measuring and statistically analyzing biological data. In information technology, biometrics usually refers to technologies for measuring and analyzing human body characteristics such as fingerprints, eye retinas and irises, voice patterns, facial patterns, and hand measurements, especially for authentication purposes.
Certificates (PKI)	A digital certificate is an electronic "credit card" that establishes your credentials when doing business or other transactions on the Web. It is issued by a certification authority (CA). It contains your name, a serial number, expiration dates, a copy of the certificate holder's public key (used for encrypting messages and digital signatures), and the digital signature of the certificate-issuing authority so that a recipient can verify that the certificate is real.
Configuration Management Repository	Configuration Management is needed for an understanding of the relationship of the components that lead to the end IT service. Configuration management is the foundation for change, problem, availability-event and IT service management, as well as disaster recovery. Specifically, configuration management enables an automated method of understanding the business impact of a change in the infrastructure. Furthermore, an understanding of specific changes to objects is recorded in the configuration management repository and used to analyze the effect and success of changes over time.
Content Filtering	Content filtering (also known as information filtering) is the use of a program to screen and exclude from access or availability Web pages or e-mail that is deemed objectionable. Content filtering usually works by specifying character strings that, if matched, indicate undesirable content that is to be screened out.
Content Monitoring	Protects against waste, fraud, and abuse by blocking access to web sites that are not work-related.
Content-Dependant Access Control	Method for controlling access of users to resources based on the content of the resource; used to protect databases containing sensitive data; for example, a patient record management system.
Continuous Availability	The highest level of availability. Designed for 100% availability including scheduled maintenance.
Data Link Encryption	Method of security by encrypting the data link layer of the transmission
Database Encryption	Method of security by encrypting the database holding specific data elements/files
Database Replication	Database replication is the periodic electronic refreshing (copying) of a database from one computer server to another so that all users in the network constantly share the same level of information.

Term	Definition
DBMS	A database management system (DBMS), sometimes just called a database manager, is a program that lets one or more computer users create and access data in a database. The DBMS manages user requests (and requests from other programs) so that users and other programs are free from having to understand where the data is physically located on storage media and, in a multi-user system, who else may also be accessing the data. In handling user requests, the DBMS ensures the integrity of the data (that is, making sure it continues to be accessible and is consistently organized as intended) and security (making sure only those with access privileges can access the data). The most typical DBMS is a relational database management system (RDBMS).
Digital Rights Management	Digital Rights Management technology allows digital content to be distributed securely on CD-ROMs, DVD-ROMs, peer-to-peer networks, enterprise networks and the Internet Digital Rights Management technology allows the protection of proprietary data by encrypting digital content and attaching usage rules to it. These rules determine the number of times the content will play, the type of devices upon which it should play, and so on.
Digital Signature	A digital signature is an electronic signature that can be used to authenticate the identity of the sender of a message or the signer of a document, and possibly to ensure that the original content of the message or document that has been sent is unchanged. Digital signatures are easily transportable, cannot be imitated by someone else, and can be automatically time-stamped. The ability to ensure that the original signed message arrived means that the sender cannot easily repudiate it later.
Directory Services	A directory is a collection of users, user passwords, and, usually, information about what network resources they can access.
Discretionary Access Control	Discretionary Access Control defines users, groups and computers that are allowed or denied access to an object based on rights granted at the discretion of the object owner.
Encryption	Encryption is the conversion of data into a form, called a ciphertext, that cannot be easily understood by unauthorized people.
Enterprise Access Management	Enterprise Access Management is the management of different user access entitlements to the wide range of applications and resources residing in the entire enterprise information network and systems. It involves 3 main components or processes: authentication, authorization and administration of the users.
Enterprise Anti-Virus	A multi-layer approach to providing anti-virus services within an enterprise. Typically consists of anti-virus software that is resident on clients, servers, and Internet e-mail gatway. Also typically includes a central facility for updating virus definitions.

Term	Definition
Enterprise User Administration	An enterprise user administration product provides a framework that enables easy, robust management of large, dynamic user populations and their privileges across heterogeneous IT infrastructures. An EUA product provides a single point of administration for the user management and access control services across multiple systems. An EUA product allows the organization to administer the resident or native access control services on all systems across the enterprise from a single point of control. Rather than replacing the existing access control services, an EUA product follows a "manager of managers" (MOM) architecture. An EUA product will offer out-of-the-box support for access control services across a range of operating systems (OSs), database management systems (DBMSs), e-mail and groupware applications, and enterprise resource planning (ERP), as well as other commercial off-the-shelf (COTS) applications.
Event Correlation	A management system that can aggregate event information from a variety of systems and security devices and conduct an analysis that looks for patterns that indicate an intrusion or other security related event.
Event Monitoring	A console for monitoring a variety of events (security, availability, etc.) across a multitude of platforms.
Fault Tolerance	Fault-tolerant describes a computer system or component designed so that, in the event that a component fails, a backup component or procedure can immediately take its place with no loss of service. Fault tolerance can be provided with software, or embedded in hardware, or provided by some combination.
File Encryption	Method of security by encrypting specific files
File Integrity Checking	A method of determining if unauthorized changes have been made to a file or database.
Firewalls	A firewall is a logical or physical discontinuity in a network to prevent unauthorized access to data or resources. A proxy server is associated with or part of a gateway server that separates the enterprise network from the outside network and a firewall server that protects the enterprise network from outside intrusion.
Hardware Token	A user held device that is used to generate a one time password based on a cryptographic algorithm.
High-Availability	Designed for near 100% availability except for scheduled maintenance.
Honey Pot	A honey pot is a computer system on the Internet that is expressly set up to attract and "trap" people who attempt to penetrate other people's computer systems. (This includes the hacker, cracker, and script kiddy.)
Host Intrusion Detection System (HIDS)	A host intrusion detection system monitors the events occurring in a computer system and analyzes them for signs of attack (or preparatory activity—host sweeps or port scans—before an attack).
Identity Management	Identity management is the creation, management and use of online, or digital, identities. Hundreds of millions of people around the world now use the Internet daily at home and at work, facing a multiplicity of corporate applications and e-business interfaces. Many such applications and interfaces require a unique user name and as a result, an individual typically possesses not one but several digital identities.

Term	Definition
Intrusion Detection System	An Intrusion detection System (IDS) is a type of security management system for computers and networks. An IDS gathers and analyzes information from various areas within a computer or a network to identify possible security breaches, which include both intrusions (attacks from outside the organization) and misuse (attacks from within the organization).
Key Management System	A system that provides the ability to securely generate, and distribute symmetric keys.
Load Balancing	Load balancing is dividing the amount of work that a computer has to do between two or more computers so that more work gets done in the same amount of time and, in general, all users get served faster. Load balancing can be implemented with hardware, software, or a combination of both. Typically, load balancing is the main reason for computer server clustering.
Malicious Code Management	Ensuring that the confidentiality, integrity or availability of the enterprise’s information assets are not breached by programs—typically viruses, worms or Trojan horses—written specifically to cause such harm
Mandatory Access Control	An access control service that enforces a security policy based on comparing (a) security labels (which indicate how sensitive or critical system resources are) with (b) security clearances (which indicate system entities are eligible to access certain resources).
Message Encryption	Method of security by encrypting a specific message
Network Intrusion Detection (NIDS)	An network intrusion detection system monitors the events occurring in a network and analyzes them for signs of attack (or preparatory activity—host sweeps or port scans—before an attack).
PKI	A PKI (public key infrastructure) enables users of a basically unsecure public network such as the Internet to securely and privately exchange data and money through the use of a public and a private cryptographic key pair that is obtained and shared through a trusted authority. The public key infrastructure provides for a digital certificate that can identify an individual or an organization and directory services that can store and, when necessary, revoke the certificates.
Redundancy	Redundant describes computer or network system components, such as fans, hard disk drives, servers, operating systems, switches, and telecommunication links that are installed to back up primary resources in case they fail. A well-known example of a redundant system is the redundant array of independent disks.
Role-based Access Control	Role-based access control (RBAC) is an approach in which users' privileges across all systems are defined primarily by roles. Each role represents a job function, and all the privileges necessary to do that job across all systems are associated with that role.
Rules-based (ACL) Access Control	In Rules Based Access Control an access control list (ACL) is used to determine which access rights each user has to a particular system object, such as a file directory or individual file.

Term	Definition
Server Clustering	Clustering is the use of multiple computers, typically PCs or UNIX workstations, multiple storage devices, and redundant interconnections, to form what appears to users as a single highly available system. Cluster computing can be used for load balancing as well as for high availability. Advocates of clustering suggest that the approach can help an enterprise achieve 99.999 availability in some cases. One of the main ideas of cluster computing is that, to the outside world, the cluster appears to be a single system. A common use of cluster computing is to load balance traffic on high-traffic Web sites.
Single Factor Authentication	Single factor authentication refers to the use of one of the three factors. for example, something you know, something you have or something you are.
Single Sign On	Single sign on is a session/user authentication process that permits a user to enter one name and password in order to access multiple applications. The single sign-on, which is requested at the initiation of the session, authenticates the user to access all the applications they have been given the rights to on the server, and eliminates future authentication prompts when the user switches applications during that particular session.
Smartcards	A smart card is a plastic card about the size of a credit card, with an embedded microchip that can be loaded with data, used for telephone calling, electronic cash payments, and other applications, and then periodically refreshed for additional use.
Storage Area Networks	A storage area network (SAN) is a high-speed special-purpose network (or subnetwork) that interconnects different kinds of data storage devices with associated data servers. Typically, a storage area network is part of the overall network of computing resources for an enterprise.
Switch/Router	A switch is a device that channels incoming data from any of multiple input ports to the specific output port that will take the data toward its intended destination. On an Ethernet local area network (LAN), a switch determines from the physical device (Media Access Control or MAC) address in each incoming message frame which output port to forward it to and out of. In a wide area packet-switched network such as the Internet, a switch determines from the IP address in each packet which output port to use for the next part of its trip to the intended destination.
Symmetric Keys	Similar to Public Key, but symmetric key cryptography involves the same key to decrypt the messages at both ends. The key is transmitted separately but still vulnerable to theft. Also know as secret keys.
Transport Encryption	Method of security by encrypting the transport layer of the transmission
Two Factor Authentication	Two factor authentication also know as strong authentication refers to the use of two of the three factors. for example, something you know, something you have or something you are.
Vulnerability Analysis	Ensuring that security exposures for each technology platform are identified in a timely and ongoing manner such that security breaches can be reduced or eliminated.
Vulnerability Assessment Services/Tool	Ensuring that security exposures for each technology platform are identified in a timely and ongoing manner such that security breaches can be reduced or eliminated.

Appendix B—Relevant Policies

Appendix B – Relevant Policies

Policy	Source
An Awareness Guide to Social Engineering	NIH
CIT Security Handbook	NIH
Glossary of NIH Incident Handling Terms	NIH
Guidance for Securing Data on Portable Systems	NIH
Guidance for Selecting Good Passwords	NIH
Incident Handling Guidelines	NIH
Limited Authorized Personal Use of NIH Information Technology (IT) Resources	NIH
NIH FTP Server Policy	NIH
NIH Guidance on File Sharing Security	NIH
NIH Guidance on Unwanted E-mail, Spam, and Chain Letters	NIH
NIH Incident Handling Procedures Text and Diagram	NIH
NIH Information Technology General Rules of Behavior	NIH
NIH NetBIOS Security Policy	NIH
NIH NIHnet/Firewall Policy	NIH
NIH SMTP Server Policy	NIH
NIH SQL Server Policy	NIH
NIH Web Server Policy	NIH
NIH Wireless Network Policy	NIH
NIH Wireless Network Security Standards	NIH
Policy on Chain Letters	NIH
Policy on Passwords	NIH
Policy on Warning Banners	NIH
Procedures for Handling Unwanted E-mails	NIH
Remote Access Policy	NIH
Sanitization Policy	NIH
Security Advice for Application Developers	NIH
Security Guidelines for NIH Remote Access Users	NIH
Security Incident Reporting Guidelines	NIH
System Security Plan Template	NIH
Active Directory	DHHS
Capital Planning and Investment Control	DHHS
Conducting Information Technology Alternatives Analysis	DHHS
DHHS AISSP Handbook	DHHS
Directory Services Using LDAP	DHHS
Domain Names	DHHS
Establishing an Incident Response Capability	DHHS
IT Security for Remote Access	DHHS
Personal Use Of Information Technology Resources	DHHS
Prevention, Detection, Removal and Reporting of Malicious Software	DHHS
Public Key Infrastructure (PKI) Certification Authority (CA)	DHHS
Usage of Persistent Cookies	DHHS
Use of Broadcast Messages, Spamming and Targeted Audiences	DHHS
Computer Security Act of 1987	Federal
Guide for Developing Security Plans for Information Technology Systems (NIST SP800-18)	Federal
OMB Circular A-130	Federal
PDD-63: Critical Infrastructure Protection	Federal

Client Contact Information

John F. Jones, Jr.
Chief IT Architect
National Institutes of Health
Telephone: +1-301-402-6759
E-mail: jonesjf@mail.nih.gov

Gartner Contact Information

Terry McKittrick
Gartner Consulting
Telephone: +1-703 226 4779
Facsimile: +1-703 226 4702
E-mail: Terry.McKittrick@gartner.com

Gartner

research consulting measurement community news