DOE G 205.1-1

Approved: 3-8-01
Review: 3-8-03
Expires: 3-8-05

DEPARTMENT OF ENERGY
CYBER SECURITY PROGRAM

CYBER SECURITY ARCHITECTURE GUIDELINES

U.S. DEPARTMENT OF ENERGY

Office of the Chief Information Officer

DISTRIBUTION:

All Departmental Elements

INITIATED BY:

Office of the Chief Information Officer

This page intentionally left blank.

CONTENTS

FIGURES

TABLES

1. INTRODUCTION

The Department of Energy (DOE), Chief Information Officer (CIO), Office of Cyber Security is responsible for developing Departmentwide cyber security policy and supporting guidance. In July 1999, the CIO issued DOE N 205.1, UNCLASSIFIED CYBER SECURITY PROGRAM, to establish the framework for the Unclassified Cyber Security Program, which is applicable to all DOE sites and contractors.[1] In concert with the draft Unclassified Cyber Security Policy, the DOE Cyber Security Architecture Guidelines were developed to provide a coherent, functional cyber security framework and identify architectural principles that should be considered when developing site-specific cyber security architectures (CSAs). Its principles and provisions are not mandatory but should be considered and used to the maximum extent possible in the development and implementation of a site's CSA. Each DOE organization, DOE Headquarters, field sites, and DOE contractors should develop a CSA based on an assessment of risks and requirements for that organization.

1.1 Purpose and Scope

This CSA provides a Departmentwide cyber security framework for the operation of existing systems and the development of future systems. This framework supports a common understanding of the design, implementation, and operations of Departmental cyber security resources and systems. The DOE information architecture[2] and DOE N 205.1 provide the policy and the CSA provides the guidance to ensure functionality and interoperability of security components in the DOE's information systems.

This CSA is focused on data networks, host systems, and applications used to store, transmit, or process unclassified information. However, the CSA also applies to computing resources used to store, transmit, or process classified information unless it conflicts with other requirements for the protection of classified information. As such, the CSA addresses all information processing and networking resources, systems, and data owned by or operated for all DOE activities.

A comprehensive cyber security program includes technical mechanisms and operational practices, procedures, and processes. This CSA is focused on the technical mechanisms and the interactions between them. When this CSA specifies a particular technical mechanism (such as a firewall), it does so to illustrate one possible means of meeting particular guidelines and/or policies required by DOE N 205.1. However, such a mechanism may not be the only means by which compliance with DOE N 205.1 can be accomplished. In general, the technical

---

[1] U.S. Department of Energy, Washington D.C., DOE N 205.1, UNCLASSIFIED CYBER SECURITY PROGRAM, initiated by the Office of the Chief Information Officer. Page 1, paragraph 1a.

[2] DOE/HR-0190, *Information Architecture, Volume I, The Foundations,* March 1995, U.S. Department of Energy, Assistant Secretary for Human Resources and Administration, Deputy Assistant Secretary for Information Management. (U.S. Government Printing Office, Superintendent of Documents, Washington, D.C.)

mechanisms and the operational practices, procedures, and processes in this CSA should be regarded as one possible implementation of DOEN 205.1. Contractor sites that implement this architecture will be generally in compliance with DOE N 205.1. When contractors choose an alternate implementation, they are expected to demonstrate that the guidelines and policies of DOE N 205.1 are met.

1.2 Security Framework

ISO/IEC Standard 15408 (Common Criteria)[3] establishes a framework and methodology for evaluating security products. Figure 1 describes security concepts and relationships adopted by this DOE CSA from the international standard. From the DOE perspective, the "owners" represent DOE, and "assets" refer to all information that is subject to the information security goal. "Threat agents" represent malicious and other human activity. Other threat agents (e.g., natural disaster) are outside the scope of this CSA. Throughout this document the terms used will have the meaning implied by the relationships shown.

The requirements listed in DOE N 205.1 address confidentiality, integrity, availability, and accountability. Each organization is responsible for the development of countermeasures that reduce risks by implementing the security policies of the owners of the information. The baseline CSA presented in Sections 2 and 3 provides a framework that will assist in the implementation of a risk-based security solution using a common architecture. DOE N 205.1 identifies specific requirements that must be addressed by each site in the preparation of a Cyber Security Program Plan (CSPP). These requirements are derived from an understanding of the threats. The threat categories addressed by this architecture include the following.

Unwanted Access and Use of DOE Networks or Hosts—Instances of this threat category may be used to obtain computing, communication, or information resources. Embarrassment to DOE and the organization may result if systems at one site are used as platforms to attack another site.

Denial of Service—Instances of this threat could diminish the ability of DOE to accomplish its mission. Denial of service occurs when a threat agent disrupts DOE systems in a manner that inhibits legitimate users from using the systems as intended. Inappropriate or excessive countermeasures to this or any other threat category could create a denial of service situation by inhibiting performance or system availability.

Information Theft—Instances of this threat category result in lost productivity resulting from the need to reproduce the information. It may also result (if the information is copied but not

[3] International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) Date: 1998-12-18, ISO/IEC/15408-1: 1999(E), *Information Technology – Security Techniques – Evaluation Criteria for IT Security – Part 1: Introduction and General Model.*

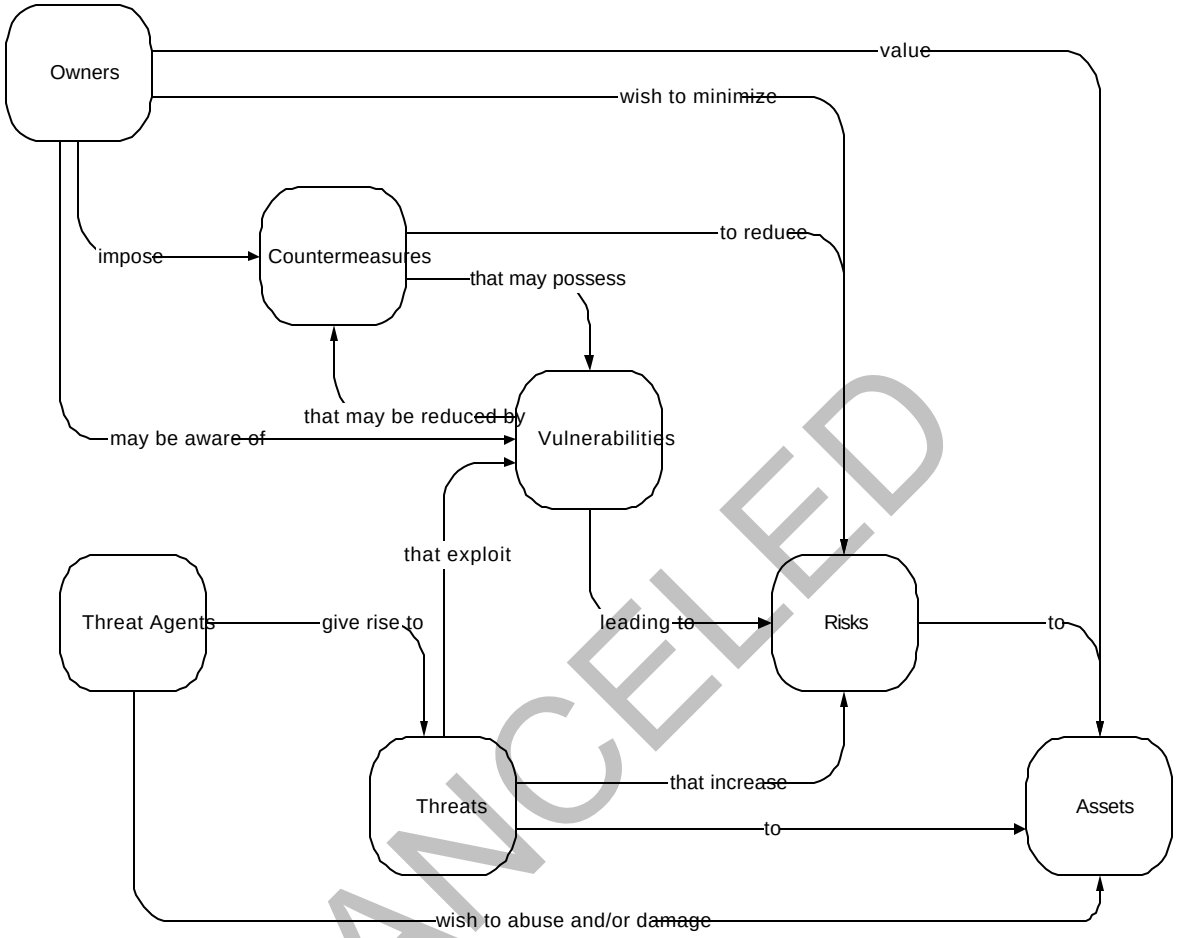

Figure 1. Security Concepts and Relationships.

destroyed) in situations in which sensitive information is in the hands of persons who may exploit it to their advantage or to the detriment of DOE, its mission, and the national security.

Malicious Modification of Information—Instances of this threat category cause errors and result in conclusions or decisions based on erroneous data. Undetected, threats in this category can create havoc for projects that rely on the integrity of stored information. Embarrassment to DOE and the organization may result if publicly viewable Web pages are modified.

This CSA addresses both external and internal threat agents. Although the emphasis is on malicious persons acting as threat agents, the model does not exclude well-intentioned, but nonetheless insecure actions of persons or systems that may also compromise the information security goal. An analysis of these threats leads to an awareness of the vulnerabilities and to appropriate countermeasures. Each site will need to perform a risk assessment/evaluation of the

cyber security threat to its specific conditions, information, missions, and environments. The results of the risk assessment combined with DOE cyber security policies provide guidance on how each site should implement its CSA in a manner that is consistent with this Departmentwide CSA.

1.3 Risk-Based Approach

This architecture incorporates a risk-based approach. Such an approach requires an analysis of the threats to information; the vulnerabilities that may exist in networks, hosts, and applications used to process that information; and the value of information and computing resources to DOE or to threat agents. Risk can be thought of as a function of threat, vulnerability, and value where risk increases as threat, vulnerability, and value increase. Each site must determine its risk based on the value of the information it is responsible for protecting, the threats, and the vulnerability of its systems. The results of this assessment should be used to determine appropriate countermeasures, while ensuring that the risk accepted by one enclave is not imposed on others.

An analysis of these threats in the DOE environment leads to an awareness of the vulnerabilities and to appropriate countermeasures. Each site will need to perform a risk assessment/evaluation of the cyber security threat to its specific conditions, information, missions, and environments. Mitigation of the site-specific risks in a form consistent with this CSA and DOE Departmentwide security guidance provides the basis for each site's CSPP. Implementation of a CSA is only one step of the risk-based approach. A prerequisite to implementation of a CSPP is the development of security policy and plans. DOE N 205.1 establishes local policy and planning requirements. Programs must be in place to test and evaluate the effectiveness of the implemented CSPP controls in order to ensure their continued effectiveness in the face of an evolving risk.

1.4 Mission Interoperability Clusters

The variety of security environments within DOE organizations requires a CSA that reflects this diversity. Accordingly, this CSA provides a set of core protection mechanisms to be augmented, as appropriate, within each environment. It also provides a structured approach using templates to identify how various environments use the protection mechanisms. DOE N 205.1 identifies five such mission areas:

- Unclassified National Security/Nuclear;
- Management, Administration, Business Operations;
- Industry and Other Government Research;
- Academic Research, Scientific Operations; and
- Open/Public/Unrestricted.

An architecture describes the system components and interactions between them. Such descriptions are intended to guide designers and implementers so that the resulting implementations satisfy the intended purpose. The core protection mechanism constructs of this

CSA are system components of this architecture that provide countermeasures to threats against DOE networks, hosts, and applications. Core protection mechanisms are identified in the body of this CSA. Appendixes A through G elaborate on how these mechanisms might be used for each of the five mission areas.

2. CYBER SECURITY ARCHITECTURE

DOE has defined a Departmentwide information architecture in which are identified a number of information architecture principles. Principle #4 states that—

> Security is designed into all architectural elements, balancing accessibility and ease of use with protection of data.[4]

This principle establishes the basis for a more detailed CSA. The architecture described in this section supports information architecture principle #4 and further identifies four of the CSA principles: Control Outside Visibility to DOE Systems, Control Access to DOE Systems, Control Interfaces Across Security Boundaries, and Monitor and Report Anomalous Activity. In addition, this CSA identifies baseline security controls related to these four principles. Implementation of the CSA may deviate from the baseline controls. However, any deviations must be consistent with Principle #4 and based on a site risk assessment. Table 1 identifies the baseline controls.

2.1 Architecture Views

The CSA describes a defense-in-depth approach. This approach is based on three views—the network view, the host platform view, and the application view. The network view focuses on network-based threats [e.g., Internet protocol (IP) spoofing, eavesdropping] and associated security controls (e.g., firewalls and routing architecture). The network view can be applied to the networks at the site boundaries and at the boundaries of networks within a site. The host platform view focuses on host-level operating system security architectures for protecting against threats such as unauthorized access or reconfiguration. The application view focuses on security for infrastructure applications [e.g., e-mail, domain name service (DNS)] and mission/corporate applications [e.g., Corporate Human Resources Information System (CHRIS) and Fossil Research Energy Database (FRED)].

The security of a particular implementation of this CSA includes mechanisms from all three views. However, considerable flexibility is possible in the selection and placement of security mechanisms within each of the three views. This defense-in-depth approach allows mechanisms

---

[4] DOE/HR-0190, *Information Architecture, Volume I, The Foundations*, March 1995, U.S. Department of Energy, Assistant Secretary for Human Resources and Administration, Deputy Assistant Secretary for Information Management. (U.S. Government Printing Office, Superintendent of Documents, Washington, D.C.) page 39.

Table 1. Cyber Security Architecture Core Principles and Baseline Controls.

Core Principle	Baseline Controls
Control Outside Visibility to DOE Systems	Each site should expose only that portion of the name/address space appropriate for external view. DNS should be configured to permit Internet name resolution for only individual host names for which the site is authoritative. Application developers should consider security requirements early in the development life cycle. Only necessary host services, as determined by the site, should be enabled. Publicly accessible information and computing resources should be designed and implemented in a manner that limits public visibility to other sites and Departmental information and computing resources.
Control Access to DOE Systems	Local implementation of the CSA will be consistent with a site-based risk assessment of the network, host, and application security. Dial-out and remote access implementations should impose security provisions consistent with those imposed upon other "on-site" users, including protection of authentication information. Source routing in IP networks is prohibited. DOE networks should only carry traffic whose source or destination address is associated with DOE requirements or that is approved by DOE. E-mail should not be configured in a manner that would allow it to be used as relay agents by unauthorized persons. Each network-attached host should be configured to require authentication before a user is granted non-anonymous access.
Control Access Across DOE Security Boundaries	Provide appropriate resistance to unauthorized changes to all components supporting connections to or from non-DOE networks, information resources, or other computing resources. Any site or enclave permitting uncontrolled entry should be treated as untrusted by other DOE sites, networks, or enclaves. Access across any DOE network boundary should be prohibited if the source address is a valid address on the destination network. E-mail traffic should be permitted only to/from on-site mail hosts whose configuration and administrative practices and procedures are controlled or reviewed by designated site cyber security personnel. Intrasite connections between campus networks with different security needs should be controlled.
Detect and Respond to Anomalous Activity	Each site will implement intrusion detection and critical event logging. Each site will implement vulnerability detection. E-mail should be configured to allow logging of message headers for inbound and outbound e-mail.

to be selected and implemented in a manner best suited to the local security requirements and usage patterns. The composite security provided by mechanisms from all three views is more important than the implementation choices in a single view.

It is unrealistic to expect to see a complete security solution from any single view. The views provide a convenient means of organizing the security components, but security is achieved only when all three views are implemented in accordance with a total CSA. An implementation of this CSA that is complete in one view, but incomplete in other views may still have significant vulnerabilities. Implemented properly, this CSA will result in a high level of confidence that DOE information is secure from unauthorized access and use.

As an example, security requirements for the academic research cluster are different from, and frequently less stringent than, the requirements for the business mission interoperability cluster. To achieve the total security requirement, a holistic view is necessary. Many of the access control mechanisms associated with the network view that are employed by the business mission interoperability cluster would be inappropriate or unnecessary for the academic research mission interoperability cluster. On the other hand the academic research mission interoperability cluster may depend more heavily on access control mechanisms such as intrusion detection and authentication.

2.1.1 Network View

Distributed computing makes computing resources available to any user linked to those facilities by a network. With this expanded computing capability, however, come security risks. Data stored and transmitted on local and wide area networks [local area networks (LANs) and wide area networks (WANs)] are vulnerable to unauthorized access, whether intentional or accidental, that could wreak havoc for DOE missions that rely on the integrity and confidentiality of their information. For that reason, security is a critical component in a distributed computing environment.

This CSA addresses these risks by defining a common view of network security components. From the network perspective, many of the security challenges that exist at one site may also exist at other sites. For example, some of the network components, such as DOE-owned wireless systems are commonly used and a family of common security solutions may be applicable. A common network security view across DOE facilitates the acquisition and implementation of common solutions.

In any network environment, information security threats originate from external as well as internal threat agents. Intrusion by external threat agents may expose additional vulnerabilities in the network or provide partial information about persons, organizations, or activities within the enterprise that, when combined with other available information, could be used to cause damage or embarrassment. Such intrusions may also result in a complete loss of confidentiality and integrity because a malicious outsider may be able to obtain control of internal hosts. Threats from outside the network require boundary protection to prevent unwanted intrusion.

The insider threat could have the same consequences as the external threat, but the countermeasures are quite different because the perpetrator has been authorized access to one or more DOE information systems with an implied level of trust to protect sensitive information. Insider threats result from maliciousness (behavior intended to cause intentional harm), negligence (risky behavior, either intentional or careless, with potentially adverse consequences), or ignorance (unwitting violation of security policy or best practices).

The DOE CSA includes components designed to mitigate each of these types of threats. However, this architecture, and in particular, this view, only addresses the security components of the DOE information architecture. It does not address training, personnel reliability programs, or other approaches (e.g., threat of criminal or administrative reprisal) to mitigate the insider threat.

In summary, when determining the network security requirements for an enclave, one must consider—

- the external threat;
- the degree to which the enclave network structure, services, and resources should be exposed to external view and/or access;
- the type of network intrusion detection and response appropriate for the enclave;
- which network services are essential for business/mission operations (e.g., file transfer, email, DNS, World Wide Web, remote access, network management, collaboration, multi-media);
- best industry practices for securing essential network services and the risk tradeoffs associated with alternatives that may provide greater access, performance, or functionality;
- the ways that enclave network resources might be exploited to cause harm to external networks/enclaves; and
- alternative controls at the host and application view that complement network controls.

2.1.2 Host View

The network view of the CSA focuses on protecting the network from unauthorized access and/or use and the protection of information in transit. However, if the individual hosts that store and process information and provide access to the network and applications are not secure, then the security architecture is ineffective. The host platform view describes a host-centric security model to complement the other views of the DOE CSA.

A host platform is the access point to the DOE network, applications, and sensitive data (e.g., data that are protected by privacy laws/regulations, proprietary, UCNI, or otherwise not intended for release beyond the protected enclave). As such, it is the ultimate target of persons

intent on gaining access to data or controlling network resources. Although denial/disruption of service is frequently an end in and of itself, gaining access to the network may be only a stepping stone to the real target, a host platform containing sensitive data. Host platform security measures also aid in dealing with the difficult potential "insider threat." Even trustworthy individuals may inadvertently compromise sensitive information if host platforms do not protect against such activity. However, from the host platform view it is the result, not the intent, which is of concern. The network provides the transportation to the target (sensitive information) and the application provides a means to view and manipulate the target. But the host platform provides the ultimate data storage repository and application services. The security mechanisms defined in the network and application/data views of this CSA must be complemented by viable host security mechanisms.

Modern operating systems such as Unix and Microsoft Windows NT include mechanisms to protect the integrity of operating system data structures. They also provide mechanisms to restrict access to the system and maintain the confidentiality of each user's data. However, it is incumbent on the system administrator to establish and maintain the security policies for such systems. The critical components of a host platform that must be protected from unauthorized use are—

- operating system data structures (e.g., the NT Registry);
- persistent storage (e.g., local hard disk);
- services (e.g., network protocols and other user-level resident applications); and
- external interfaces (e.g., SCSI port, serial/modem ports, removable storage, network adapter interface).

As illustrated in Figure 2, three layers of security protect these critical components. The outermost layer consists of vulnerability/intrusion detection mechanisms. No system that is connected to a network is 100 percent immune to latent operating system defects that may be exploited. Also, every system of policies requires some degree of monitoring and enforcement. The vulnerability/intrusion detection layer provides this system of checks and balances.

The next layer, identification/authentication, is used to positively and uniquely identify users attempting access to the system. Once positive identification has been established, the user's rights to access all or portions of local storage, external interfaces, and system services are based on the person's role and the permissions conferred to that role. The inner-most layer, the access control layer, enforces role-based access to these system components.

- Not every host system will have all of these components. For example, many systems (routers, diskless workstations) have no hard disk or other persistent storage. Hence, the architecture components and interactions described below may not be relevant in every situation. For each host platform system, managers must determine the level of risk and the appropriate degree to which each of the following CSA components must be implemented.

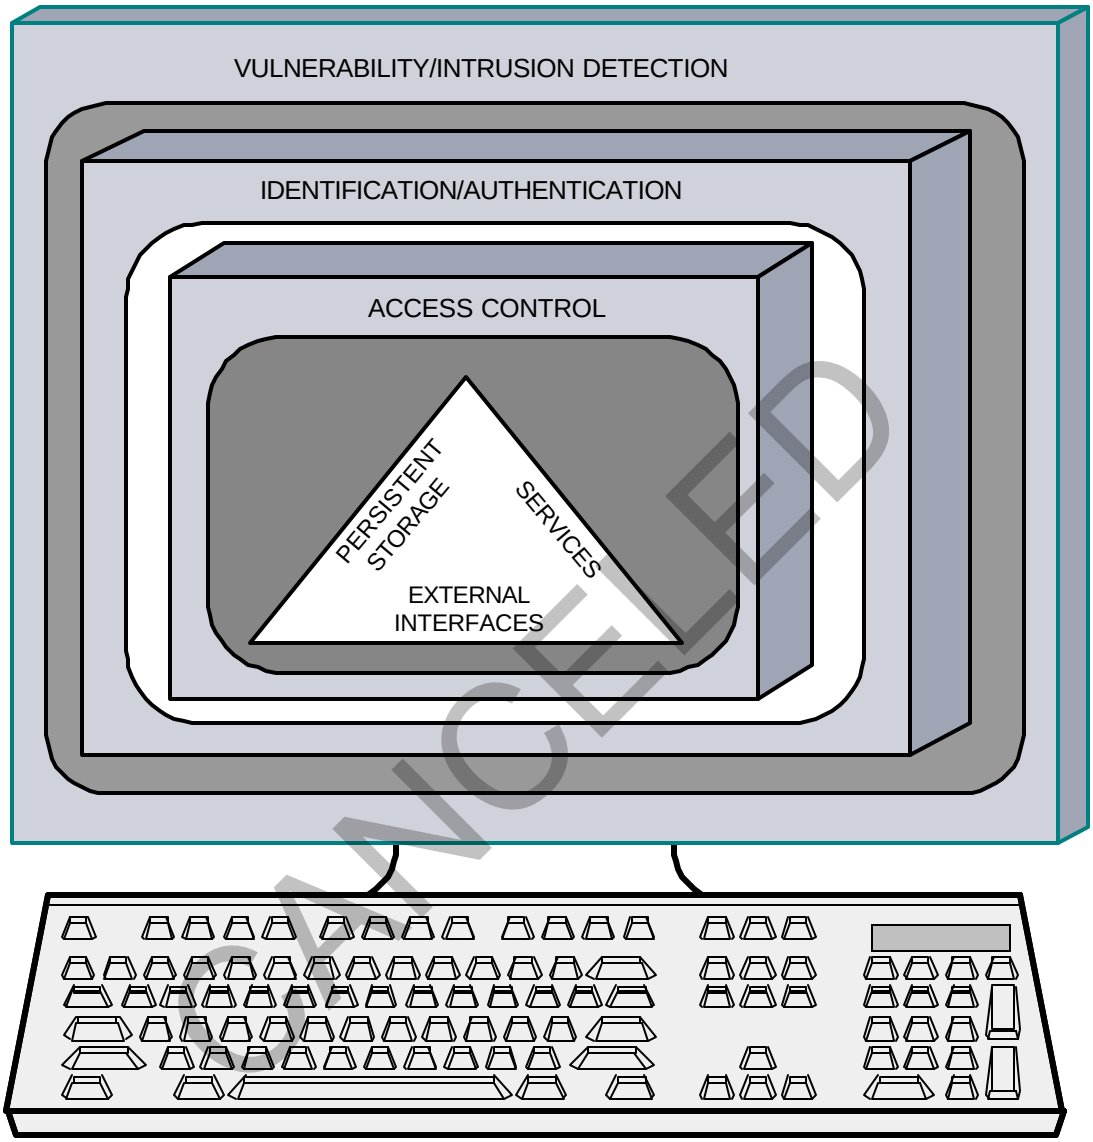

Figure 2. Host Platform Security Architecture.

In summary, when determining the host security requirements for an enclave one must consider—

- the insider threat;
- which host services are essential for business/mission operations (e.g., file transfer, e-mail, DNS, World Wide Web, remote access, network management, collaboration, multi-media);
- authentication/authorization/auditing alternatives for system and data access;

- integrity and protection of stored data;
- the type of host vulnerability/intrusion detection and response appropriate for the enclave; and
- controls at the network and application view that complement host controls.

2.1.3 Application View

A comprehensive CSA must also address the threats to applications and data. Some of these threats cannot be addressed by the services and mechanisms implemented in the network or host views. For example, applications that do not authenticate the identity of the application user can expose the application data in ways that compromise the information security. Network layer security mechanisms cannot detect and protect against application protocols that are not secure. This architecture view provides mechanisms for protecting the integrity and privacy of transmitted data as well as for authenticating the identity of application users.

The application view focuses on improving the security of enterprise applications by protecting access to the applications and data on the basis of user identity and/or role.

In summary, when determining the application security requirements for an enclave, one must consider—

- authentication/authorization/auditing alternatives for application access/use;
- integrity and protection of application data;
- application user roles and privileges;
- the location (same enclave, trusted remote enclave, other) and type of host (e.g., operating system, common use workstations) access to the application;
- the network connectivity between interacting components of the application (e.g., private versus public network, lower-layer encryption services);
- pre-existing security services/protocols/infrastructure [e.g., public key infrastructure (PKI), distributed computing security environment (DCE)];
- available security features embedded in commercial off-the-shelf (COTS) applications;
- application security frameworks for non-COTS application development; and
- alternative controls at the network and host view that complement application controls.

2.2 Security Enclaves and Protection Mechanism Types

The views presented in Section 2.1 capture the defense-in-depth approach employed by DOE for implementing cyber security. Each implementation may choose stronger mechanisms in one view than in other views to better accommodate operational requirements while addressing the security concerns. These views only provide one dimension to the defense-in-depth approach. Two other dimensions more fully describe this cyber security architecture. The first is the security enclave and the second is the type of security mechanism.

2.2.1 Security Enclaves

Each implementation of this architecture describes the protection mechanisms within the framework described in Section 2.1. However, implementers most frequently must address the specific needs of multiple security enclaves. Enclaves are defined in DOE N 205.1 as a set of information and processing capabilities that are protected as a group. The information processing capabilities may include networks, host, or applications. Typically an enclave would be located within a single implementation (e.g., site) of this architecture. However, it is conceivable that an enclave may cross network or geographical boundaries.

In addition to the mechanisms described for the network, host, and application views, the architecture templates presented in the appendixes to this document may identify specific mechanisms appropriate for implementation at an enclave level. Some mission interoperability clusters have well-defined protection requirements and implementation descriptions. Other mission interoperability clusters may have protection requirements that transcend network, host, or application boundaries. These unique requirements are not identified in the CSA description but should be documented in the implementation of the CSA.

2.2.2 Protection Mechanisms Types

Each protection mechanism presented in the core constructs and the more detailed mission area constructs presented in the appendixes may be characterized on the basis of whether the primary purpose of the mechanism is to detect, prevent, or correct a violation of the security policy. Therefore, in addition to the flexibility to provide security by selecting where to apply the mechanism (network, host, or application) and how aggressively the mechanism is applied for a given view, there is also flexibility in determining the type of mechanism (detect, prevent, or correct).

2.3 Validation

Validation is the process used to ensure that an implementation of this architecture continues to operate as intended and that such operations satisfactorily mitigate assessed risk. Although validation is largely a management and administration function, it employs certain technical mechanisms to assist the management team. Validation mechanisms may be automatic or manual. Examples include incident analysis, peer review, policy/procedure review, threat review

(manual) and vulnerability scan, and host-based intrusion detection (automatic). Each mission area template includes identification of the validation mechanisms.

3. CORE CONSTRUCTS

Section 2.1 described the three security architecture views, which included factors to consider when selecting available constructs for implementing cyber security within an enclave. This section identifies the core constructs applicable throughout DOE that can be tailored and extended to meet local objectives. Table 2 identifies the core constructs applicable to all enclaves. The remainder of this section gives a high-level description of these constructs, which may appear in an implementation of this architecture based upon the three views presented in Section 2.1.

Table 2. Core Constructs.

Network View	Host View	Application View
Boundary Protection Services	Vulnerability/Intrusion Detection	Public Key Infrastructure
Intrusion Detection Services	Identification and Authentication	Embedded Application Security

Each mission interoperability cluster elaborates the description of these constructs as appropriate in order to fully describe an implementable security architecture. Such descriptions appear in Appendixes A through E. Table 3 summarizes the protection mechanisms for each mission area enclave as described in the appendixes.

3.1 Network

The network view includes the boundary protection services and intrusion detection services constructs. The following high-level discussion of these constructs is intended to serve as a basis for the detailed presentations of each mission interoperability cluster in the appendixes.

3.1.1 Boundary Protection Services

The principal construct of this architecture relating to the network view is the boundary protection services (BPS) construct. The BPS is the component of the enterprise architecture with primary responsibility for access control to the campus network and the first-level protection of information resources for the site. In principle, the BPS will—

- selectively control the flow of traffic through it,
- selectively hide local site details,
- facilitate protection of data in transit,

Table 3. Protection Mechanism Summary.

Protection Mechanism	Network					Host					Application				
	Open/Public	Academic Research	Business	Industry/Other Research	UCNI/NNPI/Export	Open/Public	Academic Research	Business	Industry/Other Research	UCNI/NNPI/Export	Open/Public	Academic Research	Business	Industry/Other Research	UCNI/NNPI/Export
Firewall	✓	✓	✓	✓		✓	✓	✓	✓						
User Authentication	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Intrusion Detection					✓	✓	✓	✓	✓	✓					
Anti-Virus	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓					✓
Change Detection					✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Remove Unnecessary Services	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓
Audit Usage					✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Configuration Management	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Discretionary Access Control					✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Continuity of Operations						✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Encryption in Transit					✓			✓	✓	✓			✓	✓	✓
Encryption in Storage										✓					✓
Application Management													✓	✓	
Event Logs	✓		✓	✓		✓	✓	✓	✓		✓	✓	✓	✓	
Design Specification of Security Requirements					✓					✓					✓
Vulnerability Scans	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓

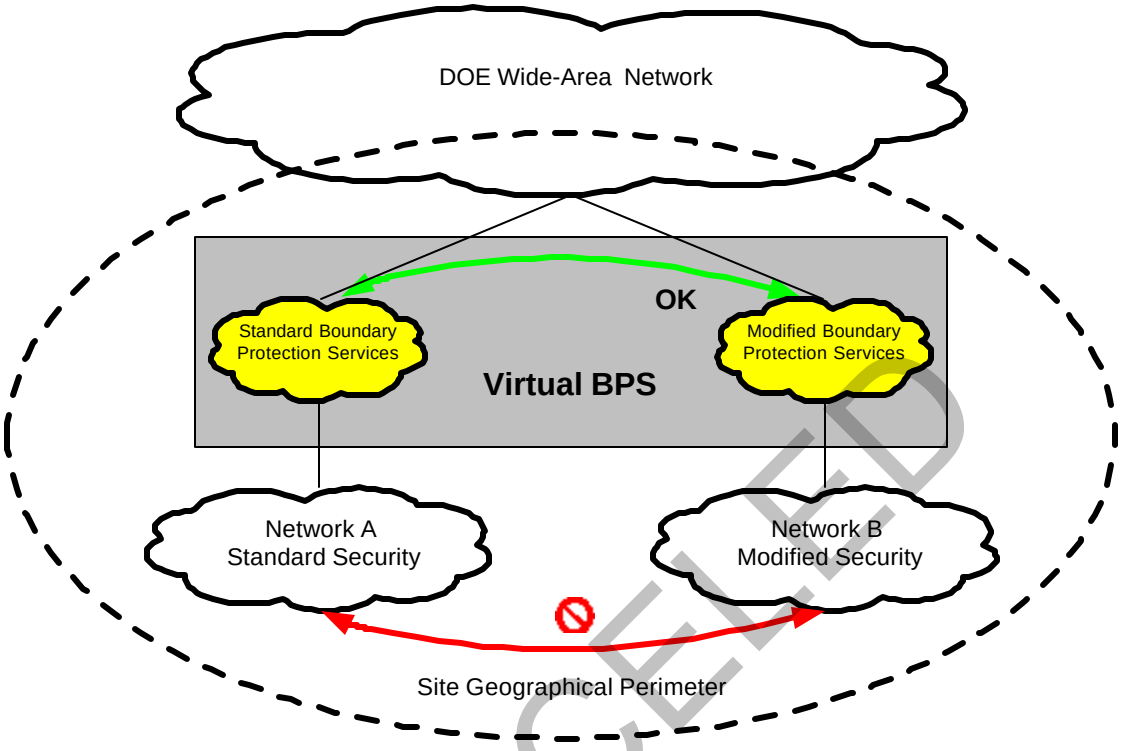

Figure 3. Virtual BPS.

- monitor network activity for anomalous behavior,
- resist unauthorized use, and
- protect itself from unauthorized change.

The degree to which each of these characteristics applies to a given BPS depends on a risk assessment conducted by and unique to the given site.

Each site may have several internal networks that require boundary protection. Each of these networks requires some baseline minimum level of information security protection, but there may be differences in the degree and type of protection required. Figure 3 shows that in such cases each network has a BPS configured to meet the perimeter protection needs of the supported network. Direct communication between two such networks almost always compromises the security posture of one or both networks and is therefore not recommended by this architecture. However, communication may be accomplished if controlled properly, usually by the BPS as indicated in the figure, or by a high-assurance “gate-guard” for specific types of information exchange between the two networks. Any entry point to a site that is not managed as a BPS

entry point should not be permitted to connect via backdoor to any site network/resource protected by a BPS. Otherwise the services provided by the BPS are undermined.

The shaded box is labeled "Virtual BPS" because it may be possible for two BPSs to share the same equipment, while performing logically as two separate BPSs, and thus provide two distinct security protection environments. It is also intended to convey that the physical placement of the BPS components may be distributed physically across the site even though they are controlled and configured as any other centrally managed BPS component on the site. The virtual BPS may be implemented as multiple physical BPSs or as a single BPS with multiple personalities, each serving to protect a different network with mechanisms appropriate to the circumstances. The key to success is that all security components are operated and managed as a security system no matter where they reside.

The BPS arrangement in Figure 3 is suitable at sites where two (or more) subnets have security needs that are different from one another but do not have a "stronger than" type of relationship. However, it is frequently the case that all subnets on a site can share some minimal subset of the total perimeter protection policy. In such cases it may be more appropriate to establish a hierarchical arrangement of BPSs or BPS components. Figure 4 illustrates this alternative. The virtual BPS for network 2 consists of the combined policies and protection mechanisms of networks 1 and 2. Although the hosts of network 1 are willing or required to accept more risk than those in networks 2 or 3, direct communication between two networks in such a hierarchy could jeopardize the security posture of the network lower in the hierarchy. The BPS serves to contain sensitive information (outbound) as well as control external access (inbound). Sound configuration management for the BPS in the architecture shown in Figure 4 is required to ensure that changes made to one BPS don't weaken the overall protection expected at a higher-level BPS.

This architecture does not specify the degree or type of protection (e.g., permitted protocols) required for the respective protection levels. Instead it simply recognizes that differences do exist and provides a framework for accommodating differing needs, even on a single site.

3.1.2 Intrusion Detection

Each implementation of this architecture will include an intrusion detection component. This component may include host-based intrusion detection in addition to network-based intrusion at the enclave boundary. Factors affecting selection and configuration of intrusion detection components include the need for—

- near real-time detection,
- automatic reporting to a central location,
- automated response/countermeasures/mitigation, and
- timely update of intrusion patterns.

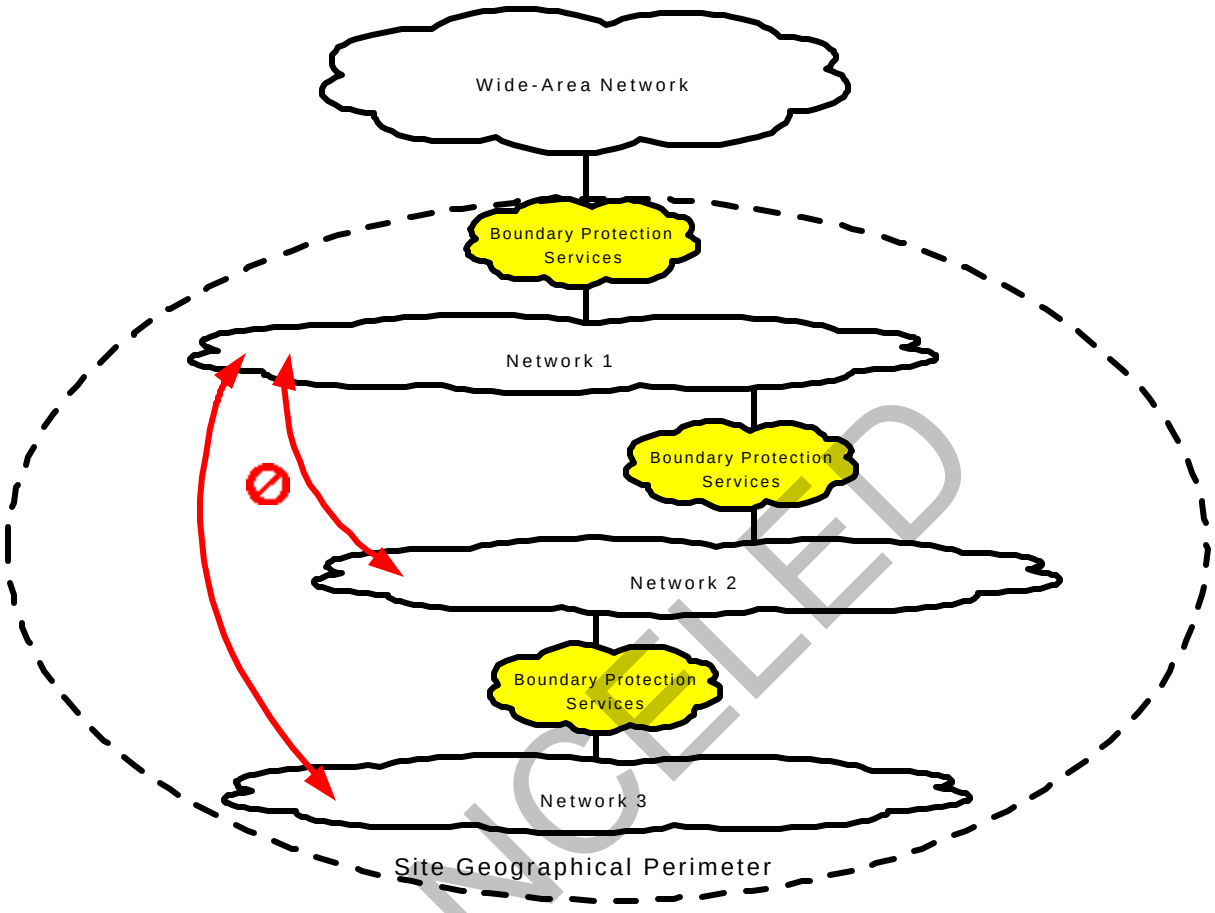

Figure 4. Hierarchical BPS Architecture.

The selection and configuration of intrusion detection depends on a risk assessment and the allocation of security mechanisms between the network, host, and application views. For example, a screened-subnet architecture for the BPS may permit the network-based intrusion detection host to more easily see both inbound and outbound traffic and therefore to correlate and analyze traffic patterns better than intrusion detection implementations in an alternative dual-homed firewall architecture. However, some mission clusters may require the added security of a dual-homed firewall for all ingress or egress. Similarly, the open/public mission interoperability cluster may depend more on the intrusion detection component than other security mechanisms in order to support open collaboration with peer labs/universities/research organizations without the more restrictive access control mechanism required by other mission clusters. The intrusion detection host is typically not visible or accessible by either the internal or external networks.

In addition to the network-based intrusion detection, the BPS may implement host-based intrusion detection with near real-time alarm functions and reporting of host platform violations of security policy.

Host-based intrusion detection agents offer capabilities not available with network-based packet-sniffing technologies. A combination of network-based and host-based technologies provides the most comprehensive intrusion detection capability. The requirements of the network view of the CSA are best met by network-based (packet-sniffing) intrusion detection systems, while the architecture requirements of the host view (Section 3.2.1) are best met by host-based intrusion detection systems. Each site will implement intrusion detection and critical event logging.

3.2 Host

The host view includes vulnerability/intrusion detection, identification/authentication, and access control constructs. The following high-level discussion of these constructs is intended to serve as a basis for the detailed presentations of each mission interoperability cluster in the appendixes.

3.2.1 Vulnerability/Intrusion Detection

Host platforms are analyzed periodically for security vulnerabilities. The method of analysis employed by DOE as part of this architecture includes both black-box and white-box techniques.

Black-box techniques use network-scanning tools to determine whether the platform is vulnerable to any network-based attacks. Black-box techniques provide no real-time security alarms. It is recommended that they occur automatically on a quarterly basis for each single-user workstation and monthly for each shared host platform, such as file/application servers or mainframe hosts. Black-box techniques are "blind" to the host platform hardware and operating system. However, the technique may "learn" this information as part of the scan and use this to discover vulnerabilities specific to the hardware and operating system of the target platform. Additionally, black-box techniques provide comprehensive coverage of the hosts on a network. New hosts that have been added to the network since the previous vulnerability analysis are discovered and subject to analysis. At a minimum, recommended black-box analysis includes IP-based vulnerabilities. However, other protocols are not excluded by this architecture.

White-box techniques use knowledge of the host platform hardware and operating system to discover vulnerabilities and intrusion attempts. White-box techniques employ agent software on the host platform to detect unauthorized use. Unlike black-box techniques, white-box techniques are capable of generating real-time alarms when—

- unauthorized activity is detected,
- security policies are violated, or
- unauthorized console access is detected.

White-box techniques may also be used like black-box techniques to periodically assess vulnerabilities and, because the agent software is specifically designed for the host operating system, white-box techniques may be better equipped to report console vulnerabilities and other platform-specific vulnerabilities. Finally, white-box techniques are better suited for host intrusion detection.

Each site will implement vulnerability detection. The DOE CSA recommends white-box techniques on shared host platforms. In some cases, COTS products for certain hardware platforms and/or operating systems may not be commercially available. In such cases, sites should establish alternative practices for vulnerability/intrusion detection. Their use on single-user platforms (e.g., user desktop systems) is not required by this architecture, but may be required for certain mission areas.

3.2.2 Identification/Authentication

This layer of the DOE host platform security architecture confirms the unique identity of the person attempting access to the system components (local storage, services, and external interfaces). After entering his/her user identification (ID) there are three ways a person can authenticate that he/she is the true owner of that user ID:

- something he/she knows (e.g., a password or personal identification number (PIN));
- something he/she has (e.g., a token or smartcard containing a PKI private key); or
- something he/she is (e.g., a fingerprint or retina scan).

One-factor authentication relies on just one of the above. Similarly, two-factor or three-factor authentication relies on two or three forms of evidence attesting to the identity of the person attempting access to the system. Based on a site risk assessment and subsequent approval within a CSPP, individual sites/managers may determine that a strong single-factor authentication (e.g., biometric related authentication) is sufficient in certain situations. This architecture also anticipates development of a DOE PKI (see Section 3.3.1) strategy. Consequently, one possible implementation of identification/authentication that would be consistent with this architecture would involve at least the second form of authentication (e.g., a smartcard containing the user's PKI private key).

Each network-attached host will be configured to require authentication before a user is granted non-anonymous access. Identification and authentication applies to both single-user and shared host platforms. Accordingly, host platform operating systems that do not support mandatory identification/authentication are not recommended by the DOE CSA. Microsoft Windows NT® and common Unix operating systems satisfy this requirement.

Identification and authentication also applies to both console and network access. Access to local storage, external interfaces, or host services must establish the identity of the person or agent requesting access. Exceptions for certain network services can be identified for each class of host platform (e.g., it is unnecessary for DNS resolvers to identify the user as part of the query to the DNS server for internal name space queries).

3.2.3 Access Control

The access control layer of the host platform view of the DOE CSA controls access to individual objects within the host platform. These objects may be individual files, applications, or folders in local storage or they may be specific host services [e.g., Lightweight Directory Access

Protocol, ftp, e-mail] or external interfaces. Host platform access control in the DOE architecture is based on the host operating system network access facilities (e.g., NT domain accounts or Unix user management facilities) or the DOE PKI if supported by the host operating system. Based on the site risk assessment the following discretionary access control mechanisms should be considered for each host or class of hosts.

Secure Logon Facility—requires users to identify themselves by entering a unique logon identifier and a password before they are allowed access to the system. The DOE PKI may be used as the basis for this facility.

Discretionary Access Control—allows an owner of a resource to determine who can access the resource and what they can do to it. The owner does this by granting access rights to a user or a group of users.

Auditing—provides the ability to detect and record important security-related events or any attempt to create, access, or delete system resources. The outermost layers, vulnerability/ intrusion detection, monitor the audit logs or directly monitor system activity and generate alerts as determined by DOE and site security policy.

Object Reuse Protection—prevents anyone from reading information written by someone else after a data structure has been released back to the operating system.

Microsoft Windows NT® and Unix operating systems support these mechanisms. However, support for the mechanisms does not ensure their use. Each user of the system must be accountable for establishing access controls to his/her information resources that are commensurate with the value of the information and others' need to know. The host platform system administrator is responsible for implementing and enforcing mechanisms to control access to the system resources; the application/data owner is responsible for implementing and enforcing mechanisms to control access to information objects he/she creates.

3.3 Application

Because each enterprise application satisfies a unique set of requirements no one construct from the application view is required for all enterprise applications. Application developers will consider security requirements early in the development life cycle to ensure that security is integrated with functionality. Application development teams should consider the constructs identified in this section.

3.3.1 PKI

DOE is developing a PKI that can be used by application developers to—

- ensure that users of systems are who they claim to be (identification and authentication);
- provide accountability that users participated in an electronic transaction (non-repudiation);
- protect against improper modifications to information (data integrity);

- prevent the disclosure of sensitive information in storage or while in transit (confidentiality); and
- enable only authorized users to access information (access control).

The DOE PKI is an infrastructure application that can be used by the other applications security constructs identified in the following sections.

3.3.2 Embedded Application Security

From a security perspective, it is convenient to view applications as either infrastructure applications or mission/corporate applications. Infrastructure applications are supporting applications that are required across all or nearly all of the DOE enterprise. These applications may be visible to the end users (e.g., e-mail) or nearly invisible (e.g., DNS). Infrastructure applications frequently must implement common security mechanisms across the enterprise to ensure interoperability and avoid vulnerabilities caused by differing implementations. Mission/corporate applications support specific DOE projects or business functions. These may include both business and science applications. Where mission/corporate applications are used across multiple sites within the enterprise, it is important to implement application security that provides the necessary security services in a manner consistent with the network and host security architecture views. Similarly, access to some applications (and data) must be restricted to authorized users within a given mission domain or area, which may also span several sites within the enterprise.

3.3.2.1 Infrastructure Applications

Infrastructure applications are typically standards-based applications for which industry has developed a security architecture tailored to the specific needs of the application. Good examples include DNS and e-mail. For such applications the DOE architecture does not have the option to alter the security architecture established by industry consensus through the standardization bodies. Instead, DOE adopts the industry standard architecture and best practices and tailors or enhances them based on the results of a risk assessment. Accepting industry standards and best practices is not sufficient to provide designers and implementers the profile or configuration guidance needed to provide interoperability and security. Such profiles and configurations must be part of the enclave security plan. Appendixes A through E provide additional specificity where significant deviation or enhancements to industry best practices are required for the mission enclaves defined in DOE N 205.1.

Event logging is an essential component of host security. It complements the three principal constructs for host security. For example, e-mail should be configured to allow logging of message headers for inbound and outbound e-mail. Without this baseline control, response to many e-mail related threats may be limited.

3.3.2.2 Mission/Corporate Applications

Mission/corporate applications are typically custom applications supporting particular mission domains or areas. A good example is CHRIS. Other examples under development or evolution include—

- Budget Execution and Formulation System (BEFS)
- Business Information Management System-Financial Management (BMIS-FM)
- Fossil Research Energy Database (FRED)
- Technical Reports Information System
- Travel Manager
- DOE Programmatic Procurement System
- Environment, Safety, and Health Reporting System

For such applications, the DOE architecture assumes that these applications must be developed using a client-server model that permits multiple client operating systems (e.g., Windows, Unix, and Mac OS). Depending on the mission area having primary sponsorship of the application, the number of clients to be supported may reduce to a single platform. Nevertheless, because mission applications must be accessible across the DOE wide area network, it is further assumed that the application architecture permits efficient implementation in a network environment (e.g., it is based on transaction processing versus file sharing). Applications and data that do not need to be accessed or shared outside of a single site may be developed or acquired using an application architecture more appropriate to the local needs. However, if the architecture described here is suitable, consideration should be given to adopting this application/data security architecture.

There are over 150 application development tools available in the marketplace. None of these tools provide good support for interoperable application/data security with applications using another tool. In selecting a security architecture for applications and data, three industry models are recommended for consideration:

- Secure Socket Layer (SSL) Web applications,
- Distributed Computing Environment (DCE) Generic Security Services Application Programming Interface (GSSAPI) enabled applications, and
- Common Data Security Architecture (CDSA) enabled applications.

Appendix F describes each of these industry models in more detail.

Any of the above industry models can support Web-based applications. Combined with an industry security model cited above Web-based applications are recommended because they

enhance interoperability with existing security mechanisms (e.g., firewalls) in the network and host views.

No single industry model suffices for all possible application needs. However, most, if not all, mission/corporate applications can be developed using one of the above security architectures as part of the overall application architecture. DOE planners and requirements analysts for mission applications need to examine both the security and functional needs of new applications and select an appropriate industry application/data security architecture from those supported by this information systems security architecture. After considering the functional and security requirements, the following additional factors should be considered in selecting the security architecture for the application.

- What security architecture is used by other mission applications in the same mission area?
- If this application is to be used by other mission areas, what is the most common security architecture employed?
- Does this mission area have applications using a security architecture that is higher in preference than the one being considered for this application?

If this application is to be used by other mission areas, is there a security architecture that is higher in preference than the one being considered for this application that is common to all mission areas/domains requiring access to this application/data?

This page intentionally left blank.

OPEN/PUBLIC/UNRESTRICTED TEMPLATE

Category: This template addresses information and systems that have been designated by the system/information owner as accessible and/or releasable to the public. The protection requirements are focused on data integrity and system availability. Integrity is required to ensure that the data is not altered, except by authorized release authority. Availability is required to ensure that the public has reasonable access to the information systems. Malicious degradation of availability must be detected and its effect minimized in a cost-effective manner.

Examples of this type of information include—

- publicly releasable research,
- public Web pages, and
- public domain databases or information repositories.

User Community: Unrestricted public access. However, a copyright notice may limit use/redistribution of the data if so designated.

Assumptions: The following assumptions were incorporated into the development of this template.

- All information protected by this template is publicly releasable.
- Administrative procedures are in place to ensure that only information that is publicly releasable is stored or processed in the open/public environment.
- Primary threats include unapproved alteration of information and use of open/public environment for unintended purposes. Additional threats include denial of service.
- The controls and protection services in this template apply to all components of information systems used to protect this information during processing, transmission, and/or storage.

Security Concerns: The security concerns for open/public information include the following:

- integrity of information and systems,
- information and/or system availability,
- malicious code,
- reputation,
- unauthorized disclosure, and
- auditability and/or accountability.

The most significant concern is integrity of information and systems as well as continuity of operations.

Table A-1. Protection of Information/Systems Accessible By or Releasable To the Public.

Protection Measure	Network	Host	Application	Notes and Comments
Firewall	✓	✓		Allow only necessary service between enclaves. Implement connectivity restrictions between this and other enclaves.
User Authentication	✓	✓	✓	Eliminate clear text reusable passwords (e.g., use encrypted passwords or credentials, one-time passwords, token cards, or smartcards) on open/public networks.
IDS (Intrusion audit, monitor)		✓		Perform at enclave boundary or on specific host(s). Implement event triggers, automatic notification, and response.
Anti-virus	✓	✓		Network-level anti-virus is becoming more important. Consider the types of servers in this enclave, how well they are protected by other measures, and whether they need additional protection.
Change detection		✓	✓	Consider products like Tripwire/SPI/COPS. Use as appropriate (tables, executables, etc.).
Remove unnecessary services	✓	✓	✓	Remove all services and applications not essential to accomplish the system mission on each system in the enclave.
Audit usage		✓	✓	Audit resource usage in accordance with internal business procedures.
Configuration management	✓	✓	✓	Maintain operating system, network, utilities, and applications patch awareness and implement patches. Routinely scan systems for vulnerabilities. Document software configuration control and procedures including installs, authorization, operations, verification, tests (especially for security patches). Ensure that the enclave network is also maintained under configuration control.

Table A-1. Protection of Information/Systems Accessible By or Releasable To the Public (continued).

Protection Measure	Network	Host	Application	Notes and Comments
Discretionary Access Controls, e.g., Authorization to Read, Modify data		✓	✓	
Continuity of Operations, Backup/Recovery		✓	✓	Consider regularly scheduled incremental and/or full backup and image restore to ensure recovery.
Encryption (in transit)				
Application Management				
Event logs	✓	✓	✓	Implement event logging and review daily. Consider triggers, automatic notification, and response.
Vulnerability Scans	✓	✓	✓	Use vulnerability scanning (e.g., ISS, COPS). Conduct scans on a regular basis (e.g., quarterly). Conduct scans on a demand basis (e.g., following an operating system or application upgrade or after patching).

This page intentionally left blank.

ACADEMIC RESEARCH, SCIENTIFIC OPERATIONS TEMPLATE

Category: This template addresses protection of academic research information and systems. The protection requirements are focused on data integrity and system availability. Integrity is required to ensure that the research data is of the highest quality at all times and that the body of data researchers depend on is reliable to the extent possible under the scientific method. System availability is required to ensure that the information and systems are available to all authorized members of a defined collaborative group, which may be geographically distributed worldwide. At the same time protection requirements for system availability must protect the researcher from premature access to research results by competing researchers/publishers. If compromised, valuable resources may be expended on faulty premises and the reputation of researcher scientist may be at stake either because data is inappropriately modified or credit for successful research is not attributed fairly.

Examples of this type of information include—

- preliminary research results,
- planned research activity, and
- pre-publication research.

User Community: The user community affected by the security controls established in this template includes all members of any private or collaborative research effort who create, edit, access, manipulate, or view research related information. These users may or may not be employees of the enterprise and may be identified as a group (e.g., all persons with approved access) or specific individuals authorized for controlled access to the information.

Assumptions: The following assumptions were incorporated into the development of this template.

- All information protected by this template is developed with the intent that it will ultimately be moved to the open/public domain at a time and under circumstances established by the researcher or collaborative body.
- Uncontrolled release of, access to, or modification of this information can have substantial and specific consequences, including legal penalties.
- The volume of information can be significant in determining protection needs.
- The unauthorized release of, or access to, a single information item or small quantities of information can constitute a serious compromise.
- The controls and protection services in this template apply to all components of information systems used to protect this information during processing, transmission, and/or storage.

Security Concerns: The security concerns for designated sensitive management, administrative, or business information include the following:

- integrity of information and systems,
- information and/or system availability,
- malicious code,
- reputation,
- unauthorized disclosure, and
- auditability and/or accountability.

Table B-1. Protection of Academic Research Information and Systems.

Protection Measure	Network	Host	Application	Notes and Comments
Firewall	✓	✓		Allow only necessary service between enclaves. Implement connectivity restrictions between this and other enclaves.
User Authentication	✓	✓	✓	Eliminate clear text reusable passwords (e.g., use encrypted passwords or credentials, one-time passwords, token cards, or smartcards) on open/public networks.
IDS (Intrusion audit, monitor)		✓		Perform at enclave boundary or on specific host(s). Implement event triggers, automatic notification, and response.
Anti-virus	✓	✓		Network-level anti-virus is becoming more important. Consider the types of servers in this enclave, how well they are protected by other measures, and whether they need additional protection.
Change detection		✓	✓	Consider products like Tripwire/SPI/COPS. Use as appropriate (tables, executables, etc.).
Remove unnecessary services	✓	✓	✓	Remove all services and applications not essential to accomplish the system mission on each system in the enclave.
Audit usage		✓	✓	Audit resource usage in accordance with internal business procedures.
Configuration management	✓	✓	✓	Maintain operating system, network, utilities, and applications patch awareness, and implement patches. Routinely scan systems for vulnerabilities. Document software configuration control and procedures including installs, authorization, operations, verification, tests (especially for security patches).

Table B-1. Protection of Academic Research Information and Systems (continued).

Protection Measure	Network	Host	Application	Notes and Comments
Configuration management *(continued)*	✓	✓	✓	Ensure that the enclave network is also maintained under configuration control.
Discretionary Access Controls, e.g., Authorization to Read, Modify data		✓	✓	
Continuity of Operations, Backup/Recovery		✓	✓	Consider regularly scheduled incremental and/or full backup and image restore to ensure recovery.
Encryption (in transit)				
Application Management				
Event logs	✓	✓	✓	Implement event logging and review daily. Consider triggers, automatic notification, and response.
Vulnerability Scans	✓	✓	✓	Use vulnerability scanning (e.g., ISS, COPS). Conduct scans on a regular basis (e.g., quarterly). Conduct scans on a demand basis (e.g., following an operating system or application upgrade or after patching).

MANAGEMENT, ADMINISTRATION, BUSINESS OPERATIONS TEMPLATE

Category: This template addresses unclassified sensitive information that has been designated by the system/information owner as requiring protection due to the sensitivity of the information or system. This is a system or information that, if compromised, could reasonably be expected to cause damage to the business interests of the organization that owns or is responsible for the information.

Examples of this type of information include—

- enterprise proprietary or trade secrets,
- competitive information;
- Privacy Act, and
- exempt from the Freedom of Information Act (FOIA).

User Community: The user community affected by the security controls established in this template includes all users who create, edit, access, manipulate, or view information that has been specifically designated as management, administrative, or business sensitive. These users may, or may not, be employees of the enterprise and may be identified as a group (e.g., all persons with approved access), or specific individuals authorized for controlled access to the information.

Assumptions: The following assumptions were incorporated into the development of this template.

- All information protected by this template is unclassified.
- Uncontrolled release of, access to, or modification of this information can have substantial and specific consequences, including legal penalties.
- There is concern about aggregation of certain categories of this information (e.g., project cost information with source of funds).
- The volume of information can be significant in determining protection needs.
- The unauthorized release of, or access to, a single information item or small quantities of information can constitute a serious compromise.
- Some information may have a maximum lifetime of sensitivity declared by the information owner.
- The information owner will declare the sensitivity of the information.
- Threats against this information include the full range of possible threats, including nation-state, state-sponsored organizations, professional and amateur hackers, corporate espionage,

arrogant and malevolent employees, and well-intentioned but careless or untrained employees.

- The controls and protection services in this template apply to all components of information systems used to protect this information during processing, transmission, and/or storage.

Security Concerns: The security concerns for designated sensitive management, administrative, or business information include the following:

- integrity of information and systems,
- information and/or system availability,
- malicious code,
- reputation,
- financial loss,
- unauthorized disclosure, and
- auditability and/or accountability.

The most significant concern is continuity of operations. This depends of the availability requirements of specific systems or information processed or stored on those systems.

Table C-1. Protection of Unclassified Sensitive Information and Systems.

Protection Measure	Network	Enclave	Host	Application	Notes and Comments
Firewall	✓	✓	✓		Multiple firewalls or other filters may be indicated where the security requirements between enclaves/hosts is different, otherwise protection can be inherited from a filter nearer the external boundary. Allow only necessary service between enclaves. Log activities in both directions across enclave boundary. Use anti-spoofing, egress filtering, and denial of service protection. Monitor packets and log appropriately. Implement connectivity restrictions between this and other enclaves.
User Authentication	✓		✓	✓	Eliminate clear text reusable passwords (e.g., use encrypted passwords or credentials, one-time passwords, token cards, or smartcards) on internal network. Use two-factor authentication for remote access (i.e., dial-up, ISDN, DSL, Internet).
IDS (Intrusion audit, monitor)		✓	✓		IDS at this enclave primarily protects against the insider threat. Perform at enclave boundary or on specific host(s). Implement event triggers, automatic notification, and response.
Anti-virus	✓	✓	✓		Network-level anti-virus is becoming more important. Consider the types of servers in this enclave, how well they are protected by other measures, and whether they need additional protection.

Table C-1. Protection of Unclassified Sensitive Information and Systems (continued).

Protection Measure	Network	Enclave	Host	Application	Notes and Comments
Change detection			✓	✓	Consider products like Tripwire/SPI/ COPS. Use as appropriate (tables, executables, etc.).
Remove unnecessary services	✓	✓	✓	✓	Remove all services and applications not essential to accomplish the system mission on each system in the enclave.
Audit usage			✓	✓	Audit resource usage in accordance with internal business procedures.
Configuration management	✓	✓	✓	✓	Maintain operating system, network, utilities, and applications patch awareness and implement patches. Routinely scan systems for vulnerabilities. Document software configuration control and procedures, including installs, authorization, operations, verification, tests (especially for security patches). Ensure that the enclave network is also maintained under configuration control.
Discretionary Access Controls, e.g., Authorization to Read, Modify data		✓	✓	✓	
Continuity of Operations, Backup/Recovery			✓	✓	Consider regularly scheduled incremental and/or full backup and image restore to ensure recovery.

Table C-1. Protection of Unclassified Sensitive Information and Systems (continued).

Protection Measure	Network	Enclave	Host	Application	Notes and Comments
Encryption (in transit)		✓	✓	✓	Identify any information that needs to be encrypted for transmission and/or encrypted for storage. The data owner is responsible specifying encryption requirements for the information. Consider default encryption (e.g., VPN) encryption between networks and/or enclaves (e.g., at different sites).
Application Management				✓	Exercise configuration management and version control of source code over its life cycle. Use life cycle development process that includes security considerations such that security is designed in, not added on. Use peer reviewed test plans with stakeholder participation and sign-off on risk. Consider all application patches and apply as appropriate.
Event logs	✓	✓	✓	✓	Implement event logging and review daily. Consider triggers, automatic notification, and response.
Vulnerability Scans	✓	✓	✓	✓	Use vulnerability scanning (e.g., ISS, COPS). Conduct scans on a regular basis (e.g., quarterly). Conduct scans on a demand basis (e.g., following an operating system or application upgrade or after patching).

This page intentionally left blank.

INDUSTRY AND OTHER GOVERNMENT RESEARCH TEMPLATE

Category: This template addresses the protection of third party unclassified systems or information that must be protected in accordance with the sponsor's requirements, as a good business practice, or to avoid compromising other systems. This includes systems or information that, if compromised, could reasonably be expected to cause damage to the business interests of the sponsoring organization that owns the system or information.

Examples of this type of information include—

- Cooperative Research and Development Agreements (CRADAs),
- third-party proprietary or trade secrets,
- competitive information, and
- any other information or systems designated by the sponsor for special protection.

User Community: The user community affected by the security controls established in this template includes all users who create, edit, access, manipulate, or view information on or associated with third-party systems and staff members or other persons designated by the third party to access and use the systems or information. These users may, or may not, be employees of the enterprise and typically will be identified as specific individuals authorized for controlled access to the systems or information.

Assumptions: The following assumptions were incorporated into the development of this template.

- All information protected by this template is unclassified.
- Uncontrolled release of, access to, or modification of this information can have substantial and specific consequences, including legal penalties.
- There is a significant concern about the disclosure of proprietary information.
- In many cases, clear ownership of proprietary information developed in the course of the project is important.
- The unauthorized release of, or access to, a single information item or small quantities of information can constitute a serious compromise.
- Most or all the information has a maximum lifetime of sensitivity specified by a legal contract and usually extending beyond the lifetime of the active contract.
- The information owner will declare the sensitivity of the information.
- Threats against this information include the full range of possible threats, including nation-state, state-sponsored organizations, professional and amateur hackers, corporate espionage,

arrogant and malevolent employees, and well-intentioned but careless or untrained employees.

- The controls and protection services in this template apply to all components of information systems used to protect this information during processing, transmission, and/or storage.

Security Concerns: The security concerns for designated sensitive management, administrative, or business information include the following:

- integrity of information and systems,
- information and/or system availability,
- malicious code,
- reputation (in particular the continuing business relationship with the sponsor),
- financial loss,
- unauthorized disclosure,
- auditability and/or accountability, and
- legal obligations.

Of particular concern is limiting access to the systems and/or information based on a strict need-to-know as governed by the sensitivity of the proprietary information potentially or actually available to users.

Table D-1. Protection of Industry and Other Government Research.

Protection Measure	Network	Enclave	Host	Application	Notes and Comments
Firewall	✓	✓	✓		Multiple firewalls or other filters may be indicated where the security requirements between enclaves/hosts is different, otherwise protection can be inherited from a filter nearer the external boundary. It may be appropriate to place additional emphasis on a host firewall to further protect these systems. Allow only necessary service between enclaves/other systems. Log activities in both directions across enclave/host boundary. Use anti-spoofing, egress filtering, and denial of service protection. Monitor packets and log appropriately. Implement connectivity restrictions between this and other enclaves (e.g., specific addresses).
User Authentication	✓		✓	✓	Eliminate clear text reusable passwords (e.g., use encrypted passwords or credentials, one-time passwords, token cards, or smartcards) on internal network. Use two-factor authentication for remote access (i.e., dial-up, ISDN, DSL, Internet).
IDS (Intrusion audit, monitor)		✓	✓		IDS at this enclave primarily protects against threats from other enclaves and/or external networks. Perform at enclave boundary or on specific host(s). Implement event triggers, automatic notification, and response.
Anti-virus	✓	✓	✓		Network-level anti-virus is becoming more important. Consider the types of servers in this enclave, how well they are protected by other measures, and whether they need additional protection.

Table D-1. Protection of Industry and Other Government Research (continued).

Protection Measure	Network	Enclave	Host	Application	Notes and Comments
Change detection			✓	✓	Consider products like Tripwire/SPI/ COPS. Use as appropriate (tables, executables, etc.).
Remove unnecessary services	✓		✓		Remove all services and applications not essential to accomplish the system mission on each system in the enclave.
Audit usage			✓	✓	Audit resource usage in accordance with internal business procedures.
Configuration management		✓	✓	✓	Maintain operating system, network, utilities, and applications patch awareness and implement patches. Routinely scan systems for vulnerabilities. Document software configuration control and procedures, including installs, authorization, operations, verification, tests (especially for security patches). Ensure that the enclave network is also maintained under configuration control.
Discretionary Access Controls, e.g., Authorization to Read, Modify data		✓	✓	✓	
Continuity of Operations, Backup/Recovery			✓	✓	Consider regularly scheduled incremental and/or full backup and image restore to ensure recovery.
Encryption (in transit)		✓	✓	✓	Identify any information that needs to be encrypted for transmission and/or encrypted for storage. The data owner is responsible specifying encryption requirements for the information. Consider default encryption (e.g., VPN) encryption between networks and/or enclaves (e.g., at different sites).

Table D-1. Protection of Industry and Other Government Research (continued).

Protection Measure	Network	Enclave	Host	Application	Notes and Comments
Application Management				✓	Exercise configuration management and version control of source code over its life cycle. Use life cycle development process that includes security considerations such that security is designed in, not added on. Use peer reviewed test plans with stakeholder participation and sign-off on risk. Consider all application patches and apply as appropriate.
Event logs	✓	✓	✓	✓	Implement event logging and review daily. Consider triggers, automatic notification, and response.
Vulnerability Scans	✓	✓	✓		Use vulnerability scanning (e.g., ISS, COPS). Conduct scans on a regular basis (e.g., quarterly). Conduct scans on a demand basis (e.g., following an operating system or application upgrade or after patching).

This page intentionally left blank.

UNCLASSIFIED NATIONAL SECURITY/NUCLEAR TEMPLATE

Category: This template addresses unclassified information that has been designated by an appropriate authority as requiring protection due to the technology being used to support the national defense or may be used against the national interests. This is information that, if compromised, could reasonably be expected to cause the decrease in the level of national or economic security relative to foreign powers.

Examples of this type of information include—

- Unclassified Controlled Nuclear Information (UCNI);
- Naval Nuclear Propulsion Information (NNPI);
- Export Controlled Information (ECI);
- Cooperative Research and Development Agreement (CRADA) information used to support the national defense; and
- Proprietary and Trade Secret information used to support the national defense.

User Community: The user community affected by the security controls established in this template includes all users who require any information that has been specifically designated as sensitive. These users are identified as specific individuals authorized (need-to-know principle) to receive the information.

Assumptions: The following assumptions were incorporated into the development of this template.

- All information protected by this template is unclassified.
- Uncontrolled release of, or access to, this information has substantial and specific consequences, including significant legal penalties.
- There is concern about aggregation of this information.
- The volume of information is not significant in determining protection needs. The unauthorized release of, or access to, a single information item or small quantities of information constitutes a compromise.
- Some information may have a maximum lifetime of sensitivity declared by the information owner.
- The information owner will declare the sensitivity of the information.
- Threats against this information include the full range of possible threats, including nation-state, state-sponsored organizations, professional and amateur hackers, corporate espionage, arrogant and malevolent employees, and well-intentioned but careless or untrained employees.

- The controls and protection services in this template apply to all components of information systems used to protect (e.g., alarm systems), store, transmit, or process this information.

Security Concerns: The security concerns for the officially designated sensitive information include the following:

- Information Confidentiality. Confidentiality concerns range from the micro-level (attorney-client privilege) to the macro level (DOE only). The specific level of confidentiality expected is information dependent and may not be based on the size of the user community accessing the information or regulations and guidelines.
- Information integrity—integrity concerns depend on the type of information, the information owner expectations, and the consequence (both legal and use impact) of unauthorized changes to the information.
- System availability
- Information availability—availability concerns depend on the type of information, the information owner expectations, and the consequence (both legal and use impact) of unavailability of the information.
- Malicious code—Malicious code, including viruses, trojan horses, and unauthorized changes to applications and operating systems have the potential to impact all of these security concerns and thereby create potentially significant legal and use impacts.
- Repudiation—The ability to effectively deny making changes to information has the potential to impact all of these security concerns and thereby create potentially significant legal and use impacts.
- Reputation—Loss of reputation for proper protection and use of officially designated sensitive information can create substantial legal and information use problems and potentially jeopardize the authority to collect, maintain, and use the information.
- Unauthorized Information Access—Unauthorized access by individuals with system access but no authorization to access the information or access by individuals not authorized to access the system has the potential to impact all of these security concerns and thereby create potentially significant legal and use impacts.
- Unauthorized Information Use or Release—Unauthorized release of information by individuals who are authorized to access and use the information, or by individuals who gain unauthorized access to the information (inside and outside).
- Software vulnerabilities (known and unknown).

Table E-1. Protection of Unclassified National Security/Nuclear Information and Systems.

Protection Measure	Network	Host	Application	Notes and Comments
User Authentication	✓	✓	✓	No clear text passwords. Only encrypted passwords or credentials, one-time passwords token cards, or smartcards. Only one-time passwords, token cards, or smartcards used for incoming access from the Internet.
IDS (Intrusion audit, monitor)	✓	✓		IDS is required at the enclave boundary, major servers, terminal and remote entry systems, optional elsewhere.
Anti-virus	✓	✓	✓	
Detecting changes in data	✓	✓	✓	
Remove unnecessary services	✓	✓		
Audit usage	✓	✓	✓	
Configuration management	✓	✓	✓	
Discretionary Access Controls, e.g., Authorization to Read, Modify data	✓	✓	✓	
Continuity of Operations, Backup/Recovery		✓	✓	The host may need special parameters, tables, or software in order to support the applications, and building the host software from vendor releases only may not provide the required functioning platform for the applications.

Table E-1. Protection of Unclassified National Security/Nuclear Information and Systems (continued).

Protection Measure	Network	Host	Application	Notes and Comments
Encryption (in transit)	✓	✓	✓	No requirement to encrypt all information within the network boundary. The user is responsible for identifying the information and encrypting it for transmission outside the boundary of the enclave. The data owner is responsible specifying the type of encryption for the information. Data owner defines the need for encryption within the boundary for need-to-know separation. Provide default encryption for encryption between enclaves (e.g., at different sites) processing this type of information. Require users to encrypt such data for transmission outside the boundary of these enclaves (e.g., via encrypted e-mail.)
Encryption in Processing and Storage		✓	✓	Risk management decision by data owner determines any additional protection needed outside the boundary, including encryption, and other mechanisms, such as VPNs.
Software vulnerability awareness, detection, response	✓	✓	✓	
Design specification of security requirements	✓	✓	✓	Documented description of security technical requirements prior to start of implementation in each view.

EMBEDDED APPLICATION SECURITY ARCHITECTURES

SSL—Secure Socket Layer (SSL) provides basic connection security between a client application and a server application. The SSL protocol has three basic properties.

- The connection is private. Encryption is used after an initial handshake to define a secret key. Symmetric cryptography is used for data encryption.
- The peer's identity can be authenticated using asymmetric (i.e., public key) cryptography.
- The integrity of the data transport is ensured through the use of message authentication codes.

SSL does not provide security for data used by or created by the application that is stored on the local machine of either the client or the server. That is, it provides secure client-server communication, but does not provide security for application data at rest. Also, SSL does not provide a structured access control infrastructure. Other components of the application architecture must provide these capabilities.

Despite the limitations of SSL, its widespread use in Web browsers makes it suitable for mission applications that can be implemented as Web applications. SSL should be avoided if the application requires fine-grained access controls unless other mechanisms such as database security mechanisms readily provide this capability and it is easily integrated with other access control mechanisms used within the mission area with primary sponsorship of the application.

GSSAPI—The Generic Security Services Applications Programming Interface (GSSAPI) is an Internet standard interface providing security functions for authentication, encryption, and digital signature. The GSSAPI calls provide a relatively high-level abstraction of these security functions; however, they do not provide mechanisms for directly accessing low-level functions, such as the cryptographic primitives available through smartcards. Applications using this API will find it easier to move between security infrastructures than if they used vendor-specific security APIs. GSSAPI implementations are available for the Distributed Computing Environment (DCE), Kerberos, and some Public Key Infrastructures, including Entrust. GSSAPI comes in two versions, and implementations may differ in the options available to the application programmer. Some code modifications may be required when porting an application between different GSSAPI implementations. Windows 2000 provides a "GSSAPI-like" interface that allows porting of Kerberos GSSAPI applications to the Windows 2000 Kerberos environment. The DCE GSSAPI[5], is a relatively mature architecture providing security

[5] Request for Comments (RFC) 2078, Network Working Group,. J. Linn, Open Vision Technologies, January 1997, Generic Security Service Application Program Interface.

(authentication, secure communication, and authorization) through a variety of services and facilities. These include—

- Registry Service—Manages objects (e.g., user accounts) in the DCE security database.
- Authentication Service—Verifies or authenticates the identity of a DCE user or service.
- Login Facility—Initializes a user's security environment (supports single sign-on for all DCE-enabled applications).
- Privilege Service—Controls access to resources by comparing the credentials with the rights to the resource, which are specified in the resource's access control list (ACL).
- Access Control List Facility—Establishes access relationships between users/services and resources.

The DCE architecture provides additional services that are not security related, but support distributed applications. These services include thread service, remote procedure call (rpc), time service, and directory service. From the perspective of the application/data view of the CSA these services are of interest to the extent they support security requirements for applications and data. For example, integrating rpc with the security service protects communication. Network data can be checked for tampering or encrypted for privacy.

The primary advantage of the DCE security architecture over the SSL security architecture is that DCE supports user privilege attributes. This enables role and capability-based access control models, in which the roles and capabilities are assigned at the point of administration of the user, rather than at the point of administration of the application or service. Additionally, the DCE application framework supports tight integration of security into the application and single-sign-on for Kerberos-aware applications. Security becomes an integral component of the application architecture.

The primary disadvantages of the DCE security architecture are that it currently does not support a public key infrastructure and does not support security for data at rest. Current DCE offerings use the Kerberos shared secret technique (private key) for authentication. However, standards groups have proposed an alternate implementation, pk_init, to the Kerberos initialization (kinit) process. Pk_init employs the services of a public key infrastructure for the initial stage of authentication. Subsequent stages, including access to other Kerberos-aware applications, remain as they have always been in DCE.

If DCE is determined to be the best industry security model for the application, there are two competing APIs for integrating Kerberos into applications. An application can call directly into the Kerberos API or it can use the GSSAPI. GSSAPI is a generic API that accesses various security providers in addition to Kerberos, but applications written to use GSSAPI are not compatible with those that use native Kerberos. Because it is a standardized API (RFC 2078), it is preferable to use GSSAPI when developing new applications. Both Windows and Unix operating systems incorporate GSSAPI support through third-party products that augment or

replace utilities provided with the operating system. Windows 2000 does not have GSSAPI, but its SSPI using Kerberos implements RFC 1964. This can allow a Unix application to use the Kerberos GSSAPI to communicate with a Windows 2000 application that is using SSPI. There is an effort at SAP AG to write a GSSAPI for Windows 2000 that will call the SSPI, thus providing Windows 2000 GSSAPI compatibility.

The GSSAPI defines an interface to cryptographically implemented strong authentication and other security services at a generic level, which is independent of any particular underlying mechanism. For example, GSSAPI-provided services can be implemented by secret-key technologies (e.g., Kerberos) or public key approaches (e.g., X.509). The DOE-sponsored GLOBUS project does this with a GSSAPI using SSL.

The GSSAPI is independent of the communications protocol suites with which it is employed, permitting use in a broad range of protocol environments. In appropriate environments, an intermediate implementation veneer that is oriented to a particular communication protocol (e.g., RPC) may be interposed between applications that call that protocol and the GSSAPI, invoking GSSAPI facilities in conjunction with that protocol's communications invocations.

The GSSAPI security context construct is independent of communications protocol association constructs, allowing a single GSSAPI implementation to be used by a variety of invoking protocol modules on behalf of those modules calling applications. GSSAPI services can also be invoked directly by applications, wholly independent of protocol associations.

GSSAPI clients are not constrained to reside within any Trusted Computing Base (TCB) perimeter defined on a system where the GSSAPI is implemented; security services are specified in a manner suitable to both intra- and extra-TCB callers.

CDSA—The Common Data Security Architecture (CDSA) is a broad suite of APIs and protocol specifications being adopted by the Open Group as a layered, portable, and integrated approach to security functionality. CDSA interfaces are defined at relatively low levels and allow for replaceable cryptographic service providers at the lowest level. An extensible Common Security Services Manager (CSSM) exists at the middle level. Applications typically interface with the CSSM or with a security abstraction layer above the CSSM (which could in fact be GSSAPI). CDSA requires that all components, except those at the application (top) layer must have an associated digitally signed credential attesting to the integrity and authenticity of the component. The integrity service ensures that the code is authentic and unaltered from the time the code was manufactured to the time the code is executed. The Security Context Management function maintains the current security context of an executing application. The security context is the profile of security parameters required by the application at any given point during the execution of the application.

CDSA includes four mandatory service modules: Trust Policy Module (TPM) Manager, Cryptographic Service Provider Module (CSP) Manager, Data Storage Library Module (DLM) Manager, and Certificate Library Module (CLM) Manager.

- TPM Manager—The TPM Manager defines a common API that allows applications to request security services that require policy review and approval as the first step in performing the operation (does person X on machine Y have permission to perform an operation that could result in modification of record Z?). Approval can be based on the identity, integrity, and authorization represented in a set of digital certificates.

- CSP Manager—The CSP Manager defines a common API for accessing all of the cryptographic service providers installed beneath it.

- DLM Manager—The DLM Manager defines an API for secure, persistent storage of certificates, certificate revocation lists, and application-specific objects. It allows applications to search and select stored data objects.

- CLM Manager—The CLM Manager defines a common API that allows applications to manipulate memory-resident certificates and certificate revocation lists. Operations include creating, signing, verifying, and extracting field values from certificates.

In addition to the mandatory service modules, certain mission areas may need additional security services not explicitly provided by the CDSA framework. CDSA provides a structure for adding a new module that defines a new category of service. The Elective Module Manager enforces this structure. An example of an elective category of security services is key recovery service.

IMPLEMENTATION EXAMPLES

Boundary Protection Service (BPS) Implementation—An implementation of the BPS that is consistent with this architecture uses dual-homed application proxy firewall technology with one interface on an "outside" subnet and one "inside" subnet. This configuration, illustrated in Figure G-1, uses virtual local area network (VLAN) technology to implement both outside and inside subnets on a single layer 2 switch[6]. Hosts shown with connections on the top of the switch are on the external VLAN, those on the bottom are on the internal VLAN. The application proxy firewall and Web proxy server are on both VLANS and must have Internet protocol (IP) forwarding turned off. It would also be consistent with this architecture if two switches were used instead of one switch with two VLANS. In either case the external and

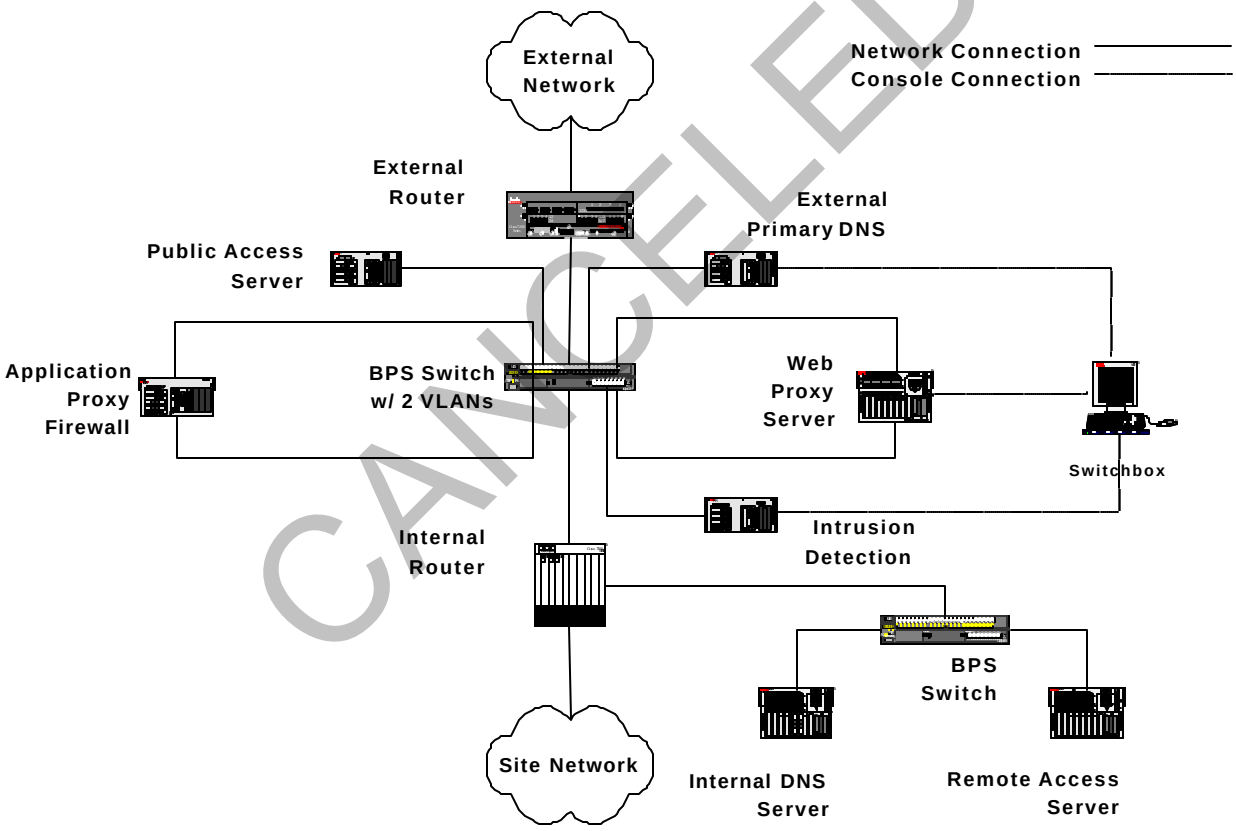

Figure G-1. BPS with VLAN Implementation.

[6] From an intrusion detection perspective, there may be an advantage to this approach because the intrusion detection host could be configured on a port that is a mirror of all traffic on both the "external" and "internal" VLANs. This might be harder to accomplish with separate switches for the external and internal LANs.

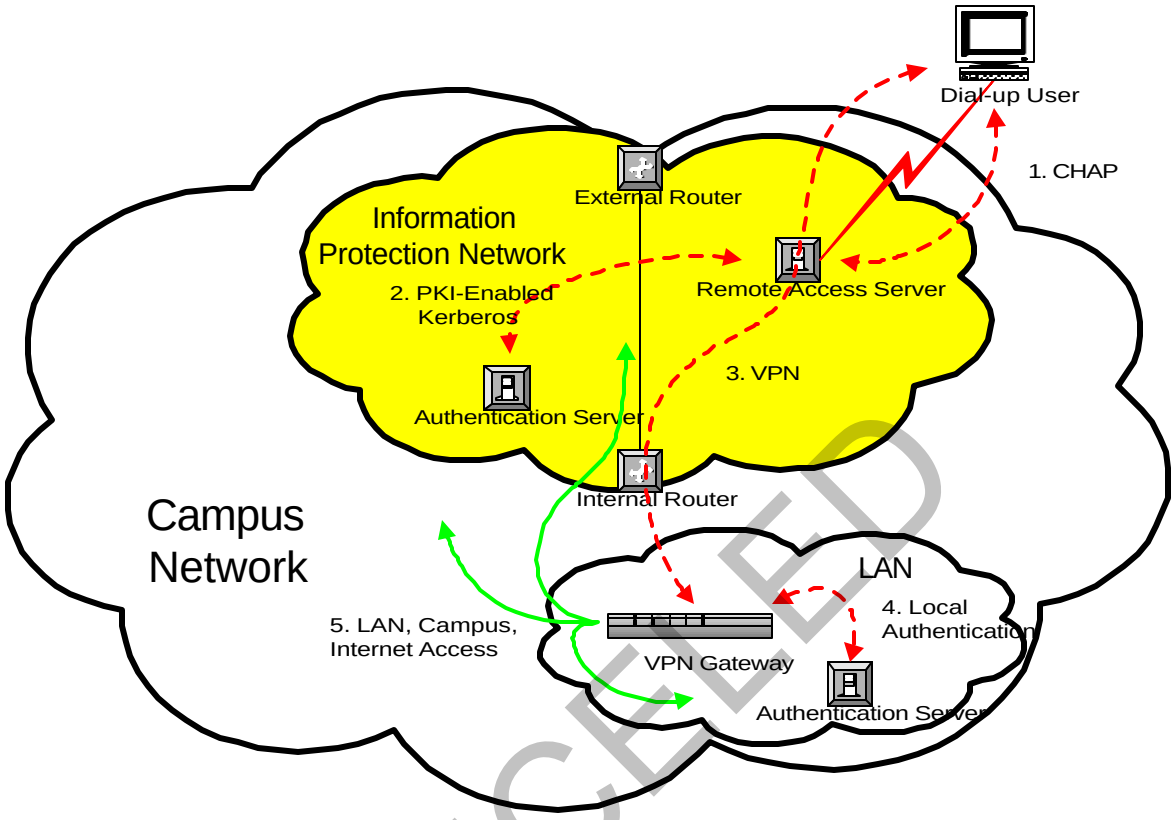

Figure G-2. Remote Access Architecture.

internal routers would typically perform access control and filtering functions to permit access to only those BPS components relevant to the external (or internal) view.

Risk can be reduced even further by avoiding the complexities of a screened subnet architecture. A simple dual-homed firewall architecture provides strict separation of external and internal traffic and may not require the same high skills and coordinated configuration management as required by screened subnet architectures. Thus dual-homed firewall implementations are also consistent with, and supported by, this architecture.

Remote Access—An implementation of remote access that is consistent with the objectives of this architecture is shown in Figure G-2. Initial implementation may employ authentication protocols other than the public key infrastructure (PKI)-enabled Kerberos shown in the figure if mature implementations of PKI-enabled Kerberos are unavailable. To preclude password sniffing, the challenge-handshake authentication protocol (CHAP) should be employed for the initial access authorization. Microsoft (MS)-CHAP or other protocols that protect against password sniffing may be used in lieu of CHAP, however protocols such as password authentication protocol (PAP) that do not provide password confidentiality should be avoided. Each LAN may establish a site-approved virtual private network (VPN) gateway for access to the LAN. Some mission traffic requirements may be able to use a common VPN gateway (not

shown in the figure) operated and maintained within the BPS rather than a local VPN gateway (shown in the figure).

This architecture does not preclude remote access from Internet workstations. Rather than a dial-up point-to-point protocol (PPP) connection to the BPS, an Internet attached user would go through the same steps, but enter through the external router. Because the architecture specifies the use of VPN technology, privacy and access control are enforced.

Support for Advanced Technology—Figure G-3 shows the security architecture for this condition/capability. For illustration purposes, asynchronous transfer mode (ATM) is shown, but the model can be applied to other technologies (including the supercomputing environment) in a similar manner. IP over ATM will be implemented in accordance with the RFC 2225[7] and the multi-protocol over ATM standard. This will permit inter-site IP over ATM capabilities for projects requiring this capability. In order to protect both the ATM network and the campus IP network, several restrictions are required. First, route servers (routing tables for IP over ATM) will advertise no route to the Internet. Doing so would circumvent the protections provided by the BPS. ATM users will be required to go through the BPS to access the Internet. This will require a traditional IP router on the ATM network with a connection to the internal router of the BPS. As shown in Figure G-3, other connections will not be permitted. Also, the internal router will block direct access to the campus IP network. If this access were allowed, other sites within the DOE enterprise would be able to circumvent the protections afforded by the BPS in violation of the principle that each site is responsible for protecting its own resources. (A similar model can be used to provide enterprisewide access to applications that are not "firewall friendly.") Because of the potential vulnerability if such interconnections are configured incorrectly, they must be treated as exceptions to usual practice and the CSPP must document safeguards to ensure the integrity of the router configuration.

The model security architecture described above for ATM can be applied analogously for other non-IP networks such as those that may exist in high-performance computing initiatives.

External Connections—Figure G-4 illustrates the constraints on external connections. First, the point of connection must be at the external router of the BPS in order to provide the necessary security protection for the site network. Generally, the access control filters between the partner network and the site should be the same as those protecting the site from any other external access. Second, the external router must be configured to protect the DOE wide area network from such connections. As shown in the figure, traffic between the site and the partner network is permitted, but traffic between the wide area network and the partner network is not permitted. Third, care must be taken to ensure that only partner network traffic is permitted onto the site network and that site network traffic is restricted to destination addresses within the partner

[7] Request for Comments (RFC) 1577, Network Working Group, M. Laubach, HP Laboratory, January 1994, "Classical IP and ARP over ATM". OBSOLETED BY RFC 2225, "Classical IP and ARP over ATM", Network Working Group, M. Laubach, Com21, Inc., J. Happern, Newbridge Networks, Inc., April 1998.

network. In other words, access to/from other networks through the partner network is not permitted, as shown in the figure.

In addition to, or instead of, an external connection to a partner network it may be appropriate to establish a shared host(s) within the BPS that is accessible only to legitimate users of the DOE network and the partner network. This would be most appropriate if the number of collaboration partners is relatively small compared to the size of the respective networks.

Domain Name Service (DNS) Access Control Example Implementation¾ Figure G-5 illustrates a set of access controls that can be implemented on site routers to restrict access to DNS servers so that only legitimate DNS activity is permitted to and from these servers.

In addition to these access controls, the external primary DNS server can be configured to use enhanced security features available in Windows NT or Unix systems as a defense-in-depth technique. Since this machine is used only for DNS resolution, only those services essential for providing this service are enabled in the operating system. Furthermore, all transmission control

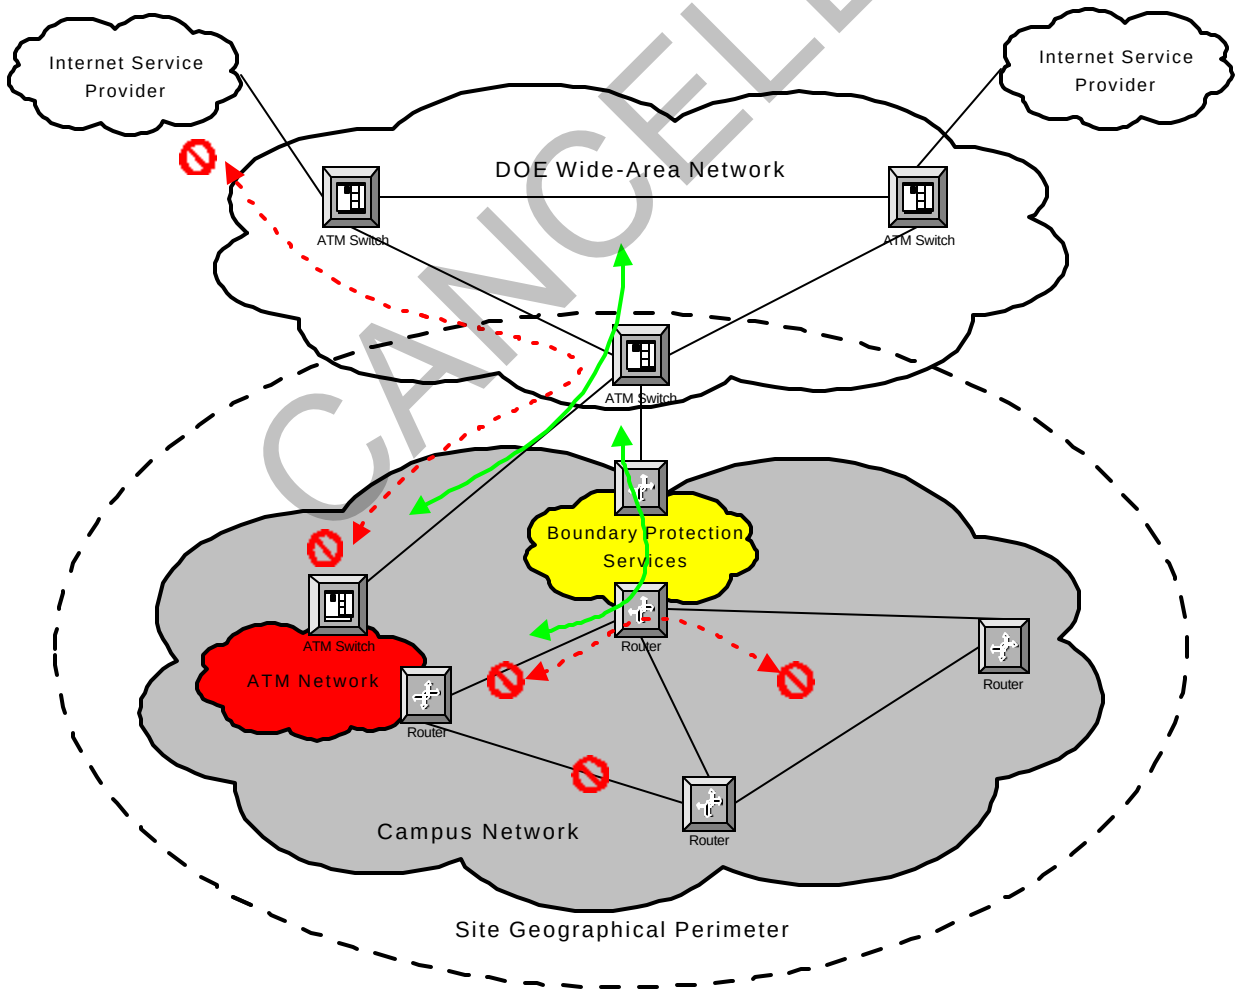

Figure G-3. Campus ATM Network.

protocol/Internet protocol (TCP/IP) ports are disabled in the operating system except those required for supporting DNS. This provides a second layer of defense if the router access controls are compromised. Additionally, only local accounts are permitted—effectively denying network login. However, if it is not feasible to administer these servers from the console (e.g., due to geographical separation between the DNS administrator and the primary DNS server) other techniques can be employed. Two such techniques are enhanced login security (encryption and strong authentication) or implementation of a stealth primary that is geographically close to the administrator and used as the source host for zone transfers by the publicly visible DNS server. DNS systems will include BIND 8 features to enhance the security of the DNS design at the application layer. The allow-query, allow-update, and allow-transfer options will be used to provide strong application layer security. They permit resolution of arbitrary Internet names by legitimate internal users (based on IP address), but restrict Internet users to queries about the zones for which the site is authoritative. Combined with mechanisms to prevent IP spoofing on site routers, these features effectively protect against cache poisoning, use of the DNS service by unauthorized users, and other DNS-related attacks.

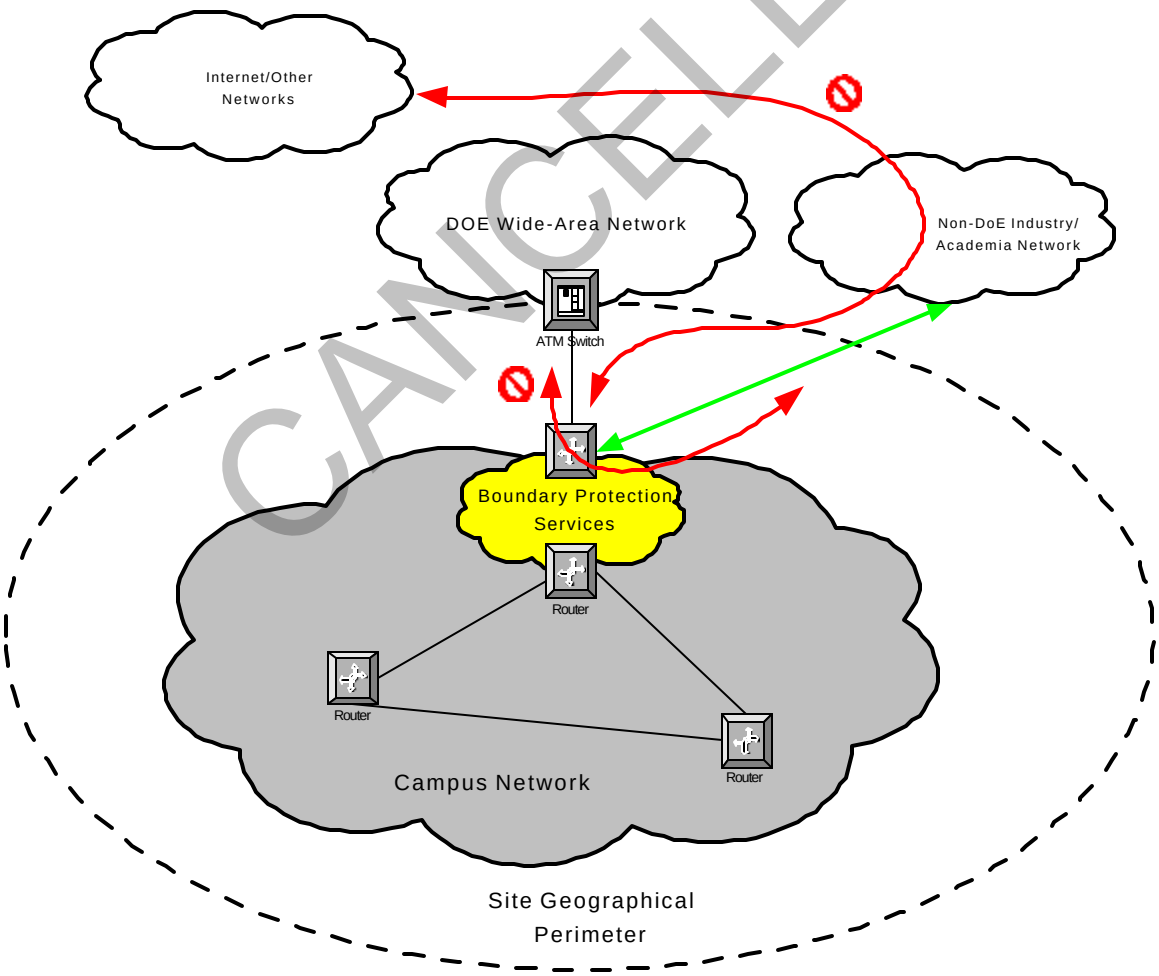

Figure G-4. External Connections.

Apply access-list 105 to inbound traffic on external router on interface facing DOE Enterprise Network

Apply access-list 106 to inbound traffic on external router on interface facing IP Protection Network

Notes					Source			Destination			
			1	access-list 105 permit	udp	any	gt 1023	host DNS Server	eq	53	
	3		1	access-list 105 permit	tcp	any	gt 1023	host DNS Server	eq	53	
		2		access-list 105 permit	udp	any	eq 53	host DNS Server	gt	1023	
		2		access-list 105 permit	tcp	any	eq 53	host DNS Server	gt	1023	established
		2	1	access-list 105 permit	udp	any	eq 53	host DNS Server	eq	53	
4				access-list 105 permit	icmp	any		host DNS Server			echo-reply

Notes

1 legitimate queries from the internet
2 legitimate reply to query from external DNS server
3 zone transfer request from the internet (External DNS server only responds to authorized secondaries)
4 permit replies to ping (optional)

Notes					Source			Destination			
			5	access-list 106 permit	udp	host DNS Server	eq 53	any	gt	1023	
	7		5	access-list 106 permit	tcp	host DNS Server	eq 53	any	gt	1023	established
		6		access-list 106 permit	udp	host DNS Server	gt 1023	any	eq	53	
		6		access-list 106 permit	tcp	host DNS Server	gt 1023	any	eq	53	
			5	access-list 106 permit	udp	host DNS Server	eq 53	any	eq	53	
8				access-list 106 permit	icmp	host DNS Server		any			echo

Notes

5 legitimate replies from external DNS server
6 legitimate queries from external DNS server
7 zone transfer request to internet (External DNS server will only send requests to authorized primaries)
8 permit pings from the External DNS server (optional)

Figure G-5. DNS Access Control List.