Network Applications Consortium

Enterprise Security Architecture

A Framework and Template for Policy-Driven Security

About NAC

Founded in 1990, the Network Applications Consortium (NAC) is a strategic end-user organization whose vision is to improve the interoperability and manageability of business-critical applications being developed for the heterogeneous, virtual-enterprise computing environment.

NAC's mission is to promote member collaboration and influence the strategic direction of vendors developing virtual-enterprise application and infrastructure technologies. NAC represents combined revenues of over $750 billion dollars, more than 50,000 network servers, and over 1 million workstations.

NAC membership radically improves the delivery of agile IT infrastructure in support of business objectives. NAC members consolidate, clarify, and communicate infrastructure technology needs to influence the IT industry and drive the evolution of standards and products.

NAC members include:

ABN AMRO
Agilent Technologies, Inc.
Bechtel Corporation
Boeing Company
ChevronTexaco
Cisco
GlaxoSmithKline
Idaho National Engineering & Environmental Lab
Johnson Controls, Inc.
Knights of Columbus
Lawrence Livermore National Laboratory
Pacific Gas & Electric Company
PricewaterhouseCoopers
Principal Financial Group
Progress Energy
State Farm Insurance
TD Bank Financial Group
The Phoenix Companies
Unisys
University of Wisconsin-Madison
Walt Disney Company

This paper is the result of NAC's Strategic Interest Group (SIG) process, a collaborative effort of a subset of NAC members whose mission is to provide a cohesive NAC viewpoint on a particular industry sector or technical topic. The following NAC members were instrumental in writing this paper:

Bechtel Corporation	*Fred Wettling*
Boeing Company	*Mike Beach*
GlaxoSmithKline	*Joe Caruso*
Idaho National Engineering & Environmental Lab	*Barry Stevenson*
Principal Financial Group	*Kevin Kelley*
Progress Energy	*Merl Ferguson*
State Farm Insurance	*Karl Hedding, Bruce Lane*
TD Bank Financial Group	*Andrew Marshall, Jim Weaver*
University of Wisconsin-Madison	*Stefan Wahe*
SAWG Core Team Co Leaders	*Mike Beach, Stefan Wahe*
SAWG Core Team Members	*Merl Ferguson, Kevin Kelley, Bruce Lane*
Project Manager & Technical Writer	*Harold Albrecht*

We welcome your feedback about this paper. For more information contact:

Doug Obeid, Chief Executive Officer
Network Applications Consortium
(808) 874-8408 or (415) 282-8670
dobeid@netapps.org
http://www.netapps.org

Table of Contents

Table of Figures

Executive Overview

> "*The information technology revolution has changed the way business is transacted, government operates, and national defense is conducted. Those three functions now depend on an interdependent network of critical information infrastructures. The protection program authorized by this order shall consist of continuous efforts to secure information systems for critical infrastructure, including emergency preparedness communications and the physical assets that support such systems. Protection of these systems is essential to the telecommunications, energy, financial services, manufacturing, water, transportation, health care, and emergency services sectors.*"[1]

Information systems security has never been more critical in this country and around the world. At the same time, protection of information systems is increasingly complex. Demand for new and improved services in both the public and private sectors is intense, and as enterprises reinvent their services infrastructure to meet this demand, traditional boundaries are disappearing. The cyber security threats lurking outside those traditional boundaries are real and well documented. Security by exclusion is more necessary and more difficult, yet not sufficient. The enterprise must also practice security by inclusion, to allow access to the services that citizens, customers, suppliers, and business partners are demanding; to allow employees and independent agents to work effectively from home; or to support some other variation on user access to the services of the enterprise.

Late in 2003 a group of NAC members began meeting the challenge of describing a common framework that would speed the process of developing enterprise security architectures for this complex environment and create the governance foundation for sustaining it into the future. How does one simplify the process of governing security by exclusion (keeping the bad guys out) and security by inclusion (allowing, and encouraging legitimate users to come in)? NAC's premise[2] is that policy-driven security architecture is essential in order to simplify management of this increasingly complex environment. As the *Corporate Governance Task Force Report*[3] states, "The road to information security goes through corporate governance." At the heart of governance are policy definition, implementation, and enforcement. To simplify security management, there must be a direct linkage between governance and the security architecture itself—in other words, policy-driven security architecture.

What is policy-driven security architecture? It starts with a policy framework for identifying guiding security principles; authorizing their enforcement in specific control domains through a set of policies; and implementing the policies through technical standards, guidelines, and procedures. It continues with a policy-driven technical framework for creating electronic representations of the policy standards,

[1] From the *Executive Order on Critical Infrastructure Protection,* George W. Bush, The White House, October 16, 2001

[2] This premise is not uniquely the NAC's. It is shared by many others in the industry.

[3] The full report is available at http://www.cyberpartnership.org/init-governance.html.

storing them in central policy repositories, and referencing them at runtime to make and enforce policy decisions. Finally, it provides a policy-driven security operations framework for ensuring that the technology as deployed both conforms to policy and enforces policy across the environment.

What is NAC's approach to designing policy-driven security architecture? It starts with defining an enterprise security program framework that places security program management in the larger context. It continues with in-depth focus on the three major components that make up enterprise security architecture: governance, technology architecture, and operations.

For governance, the NAC approach establishes the overall process, defines the policy framework that is at the heart of governance, and provides templates for security principles and policies. The principles template is derived from the National Institute of Standards and Technology (NIST) Engineering Principles for IT Security, supplemented by principles from NAC member organizations and others. The policy template is adopted directly from ISO 17799, *Code of Practice for Information Security Management*, which is an international standard that is gaining traction in the enterprise security space.

For technology architecture, the approach defines a generic framework for the management of policy-driven security services and then utilizes the framework as the basis of an overall conceptual architecture for implementing policy-driven security services. The NAC framework is based in part on the Burton Group's Virtual Enterprise Network (VEN) Security Model[4], and in part on current and evolving standards in the policy management space. It extends the policy-driven concepts beyond access management to include configuration of other security services such as border protection, cryptography, content management, and auditing. The standards are starting to gain acceptance, and their success is critical to the implementation of general-purpose policy-driven security architecture. Without these standards, centralization of policy and interoperability of the many products in the federated environment will not be possible. *Because the critical standards and implementing products are immature and incomplete, policy-driven security architecture must be considered a work in progress.* The complete solution can't be purchased off the shelf today; however, it is being implemented incrementally by vendors and users, including a few NAC members.

In addition to the overall conceptual architecture, two of the identified security services—identity management and border protection—are analyzed further to the level of service-specific conceptual and logical architecture. These two examples illustrate the logical decomposition of high-level services to the level of detail required to implement the architecture. For other identified security services, there is a template of high-level service definitions but no additional detailed perspective.

[4] The Burton Group's VEN Security Model is described in *Securing the Virtual Enterprise Network: Layered Defenses, Coordinated Policies V2*, dated May 23, 2003.

The NAC approach to security operations is to define the operational processes required to support a policy-driven security environment. These processes are of two types. One includes the administration, compliance, and vulnerability management processes required to ensure that the technology as deployed conforms to policy and provides adequate protection to control the level of risk to the environment. The other category includes the administration and event and incident management processes required to enforce policy within the environment. These operational processes are defined at a high level, not at the level of detail provided for governance and technology architecture.

Having developed the major components of NAC's enterprise security architecture (ESA), this document goes on to describe the vision, technical model, and roadmap for achieving automated definition, instantiation, and enforcement of security policy. It begins by defining the policy layers and policy automation vision:

- Start with the high-level definition of a business policy.
- Map that to a set of ISO 17799 security policies.
- Translate the security policies to detailed technical standards.
- Instantiate an electronic representation of those standards.
- Then use that representation to drive the automated decision-making and enforcement process.

The document then describes a technical model for implementing the vision and finally establishes a roadmap of user and industry actions required to enable that technical model. NAC's intent is to use this portion of the document, in conjunction with members and alliance partners, as a catalyst to organize and drive the required industry standards and technology actions. This work is starting in conjunction with the DMTF at a joint collaboration session following the NAC 2004 Fall Conference.

This document concludes with recommendations to NAC member organizations as well as vendors and standards organizations in the policy-driven security space. The key recommendation to members is that they proceed with implementation of the ESA policy and technology frameworks, recognizing that for now they must rely on manual mapping of business policies to the detailed technical standards required for decision making and enforcement, which can be partially but not fully automated. This will position members for incremental transition to policy automation products as business drivers and technology warrant—in some areas this can start today. Vendors and standards organizations are encouraged to adopt ESA as a common vocabulary; support current and emerging standards related to policy-driven security; and consider the opportunities for open, standards-based products that support a common policy automation vision.

Introduction

There is general agreement among certified security professionals and others that the overall objective of information security is to preserve the *availability*, *integrity*, and *confidentiality* of an organization's information. Effective IT security management also calls for providing *accountability* and *assurance*. Enterprise security architecture is the component of the overall enterprise architecture designed specifically to fulfill these objectives. A critical element of enterprise information security is physical security, which is the linchpin of a secure environment.

Enterprise security architecture may also be thought of as the overall framework for fulfilling these objectives while satisfying the security demands placed on the IT service organization by its customers. It includes all aspects of security governance, security technology architecture, and security operations required to protect the information technology assets of the enterprise.

The objective of this document is twofold:

- Provide a framework that serves as a common reference for describing enterprise security architecture and technology both within and between organizations.
- Provide a template that allows user organizations to select the elements of enterprise security architecture they require and to tailor them to their needs.

General Description of an Enterprise Security Program

NAC's enterprise security architecture must be understood in the larger corporate context, where it is part of an overall enterprise security program, as shown in Figure 1. It must relate appropriately to the corporate risk management, corporate IT

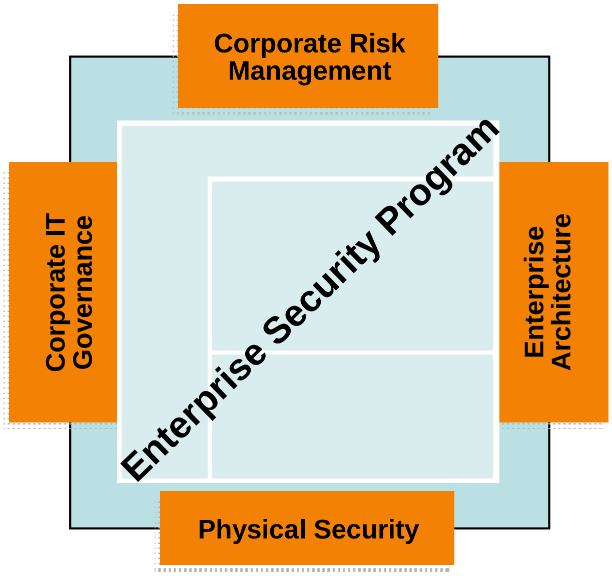

Figure 1. Corporate Enterprise Security Context

governance, enterprise architecture, and physical security programs of the enterprise. The specifics of how it relates may vary from one organization to another.

The overall enterprise security program is expanded in Figure 2 as four concentric rings of responsibility:

- Overall program management responsibility lives in the outer ring.
- Security governance responsibility lives in the second ring.
- Security technology architecture responsibility lives in the third ring.
- Security operations responsibility lives in the inner ring.

Each ring identifies key components and processes that fall within that responsibility domain. Viewed in the context of a constraints-based methodology, the components of each ring represent deliverables that further narrow the definition of what must be provided by the inner rings. Thus the requirements, strategy, planning roadmaps, and risk management assessments from the outer ring narrow the definition of what must be provided in the governance and technology architecture rings. For example, a new privacy requirement may dictate the definition of new governing principles, policies, and standards as well as the implementation of new technology architecture. The

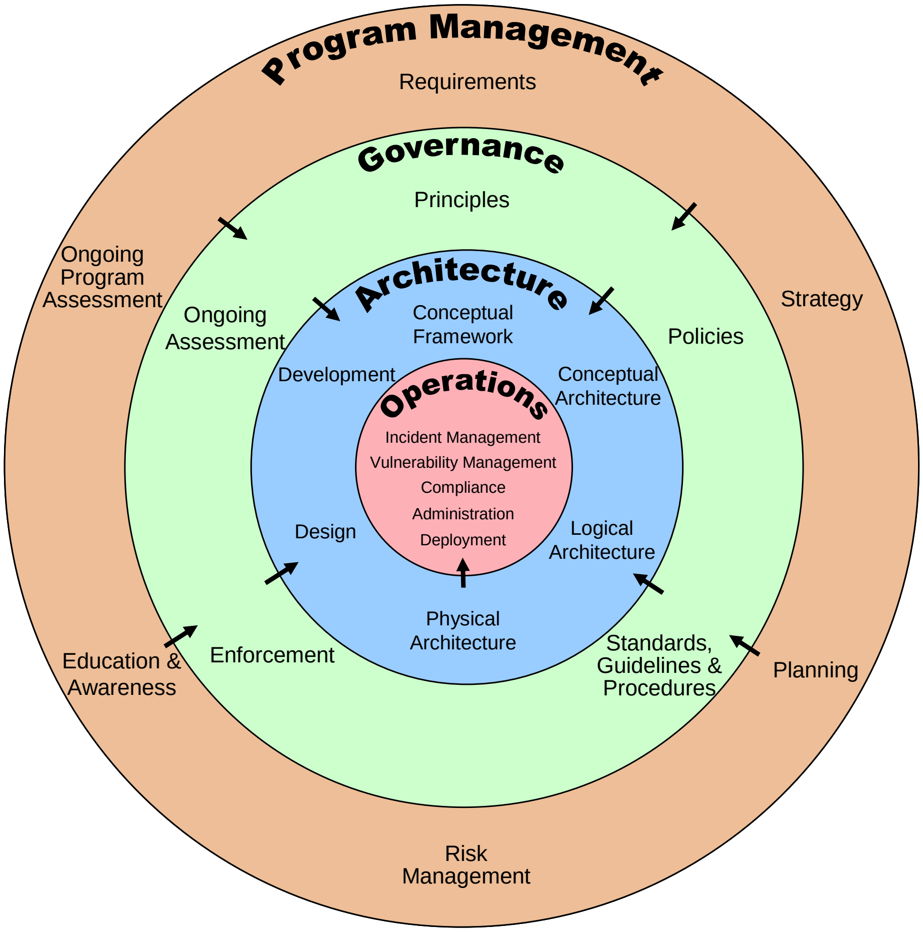

Figure 2. Enterprise Security Program Model

implementation of new standards and new architecture may in turn dictate the creation of new security processes or other capabilities within operations.

The program management functions identified in the outer ring of the enterprise security program model are considered outside the main scope of NAC's security architecture focus. In the next major section, the document focus will shift to the enterprise security architecture (ESA) components identified in the inner rings: security governance, security technology architecture, and security operations. First, however—in recognition of the importance of the program management functions—the following section describes an overall enterprise security program framework. The goal is to provide a more complete overview of the security drivers and the program management functions, and also to provide a preview of the ESA structure and show how it relates to program management.

Enterprise Security Program Framework

Figure 3 provides a more complete framework view of the enterprise security program. *Note that the rectangular boxes represent components or deliverables, while the octagonal boxes represent processes.* The framework starts with the four security drivers shown at the top, which identify the primary sources of security requirements that must be addressed. The key sources of internal requirements are the business areas, which have service-level business requirements they must meet to serve their current customers and to take advantage of new business opportunities. External requirements include security threats and legal and regulatory compliance requirements. Privacy and confidentiality are key examples of functional

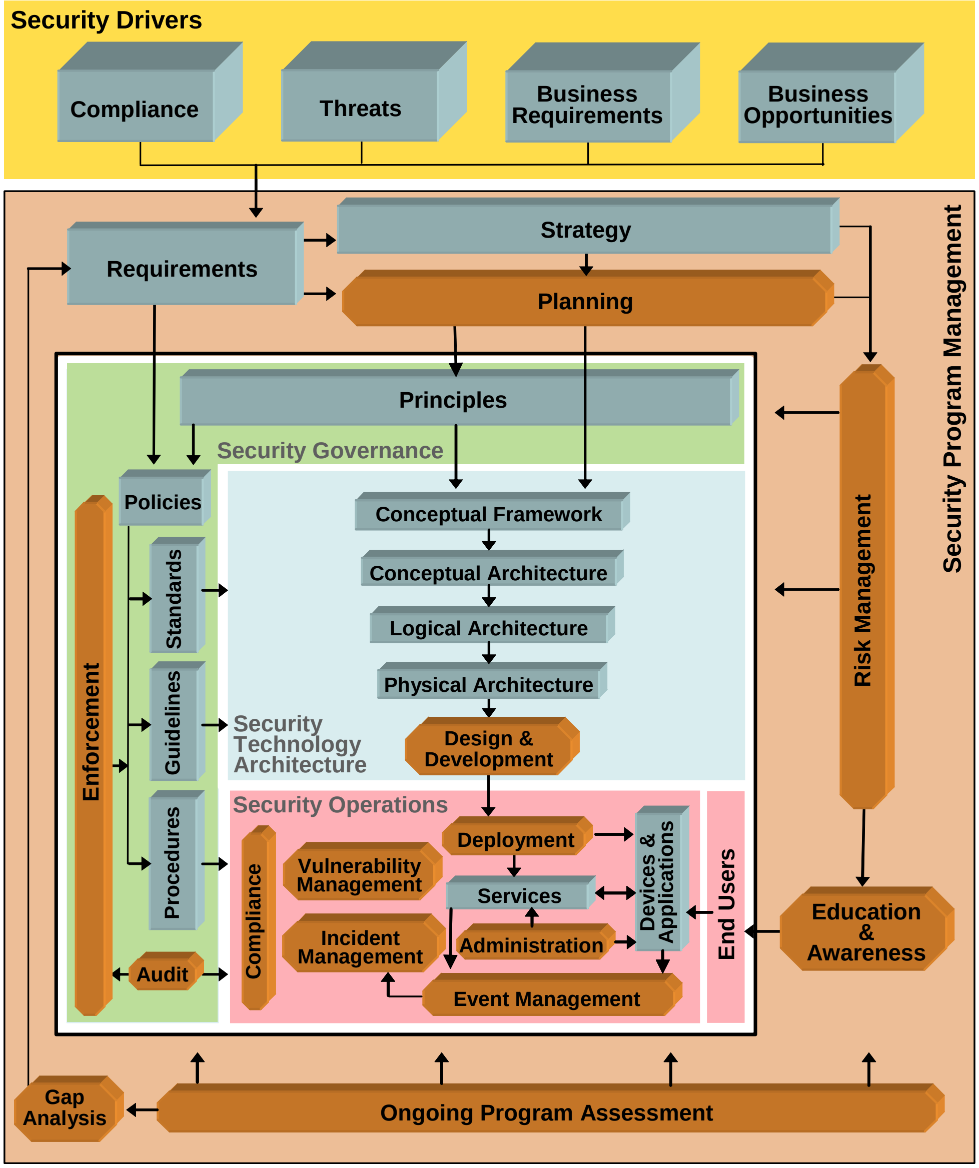

Figure 3. Enterprise Security Program Framework

requirements driven by legal requirements. Risk management may also be affected for business areas within the purview of external regulatory commissions.

Requirements drive the development of the security program strategy deliverables as well as the planning process. Risk management is the crucial process of determining the acceptable level of security risk at various points in the enterprise IT system and implementing the optimal level of management and technical control; too little control may result in financial exposure, and too much may result in unnecessary cost. Education and awareness processes are critical to the success of any security program. Ongoing program assessment and gap analysis processes provide continual requirements feedback.

The functions of the ESA components and processes are summarized below and will be described further in the subsequent sections of the document.

Governance

- Principles: basic assumptions and beliefs providing overall security guidance.
- Policies: the security rules that apply in various control domains.
- Standards, guidelines, and procedures: the implementation of the policies through technical requirements, recommended practices, and instructions.
- Audit: the process of reviewing security activities for policy compliance.
- Enforcement: the processes for ensuring compliance with the policies.

Technology Architecture

- Conceptual framework: generic framework for policy-based management of security services.
- Conceptual architecture: conceptual structure for management of decision making and policy enforcement across a broad set of security services.
- Logical architecture: provides more detail on the logical components necessary to provide each security service.
- Physical architecture: identifies specific products, showing where they are located and how they are connected to deliver the necessary functionality, performance, and reliability.
- Design/development: guides, templates, tools, reusable libraries, and code samples to aid in the effective utilization and integration of applications into the ESA environment.

Security Operations

- Deployment: assumed to be the normal IT deployment process, not a security operations process.
- Services: the core security functions defined by the security technology architecture that support devices and applications, as well as other security operations processes.

- Devices and applications: devices and applications that use ESA services and are supported by the security operations processes.
- Administration: the process for securing the organization’s operational digital assets against accidental or unauthorized modification or disclosure.
- Event management: the process for day-to-day management of the security-related events generated by a variety of devices across the operational environment, including security, network, storage, and host devices.
- Incident management: the process for responding to security-related events that indicate a violation or imminent threat of violation of security policy (i.e., the organization is under attack or has suffered a loss).
- Vulnerability management: the process for identifying high-risk infrastructure components, assessing their vulnerabilities, and taking the appropriate actions to control the level of risk to the operational environment.
- Compliance: the process for ensuring that the deployed technology conforms to the organization’s policies, procedures, and architecture.

Enterprise Security Architecture

With the enterprise security program framework as background, the focus for the remainder of the document shifts to the ESA components. As shown in Figure 4, the security program management functions now assume a background role and become part of the larger corporate context, as the focus shifts to security governance, security technology architecture, and security operations. Our goal is to describe an

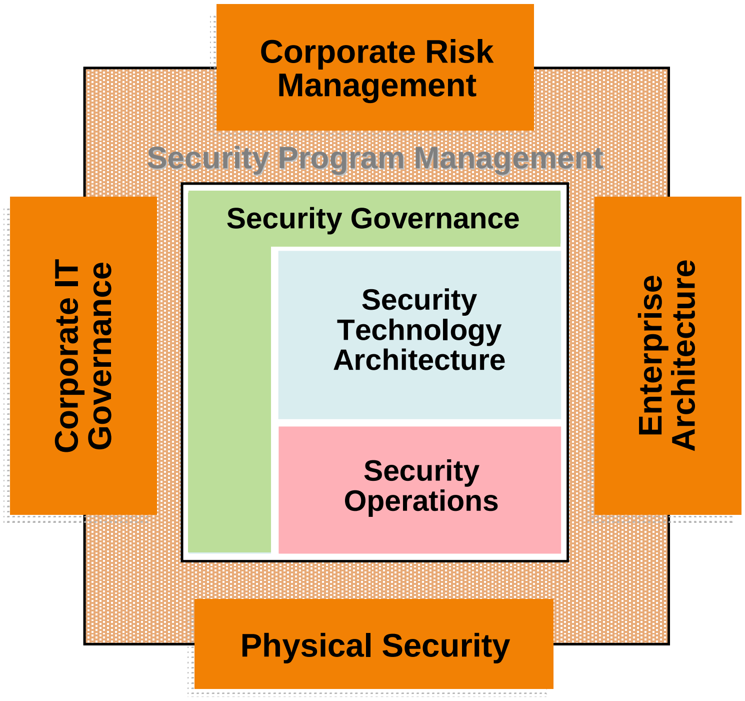

Figure 4. Enterprise Security Architecture Components

ESA framework and templates that user organizations can understand, tailor to their needs, and use as a starting point for an ESA implementation.

To effectively design and implement ESA, one needs to understand the purpose and relationships of the ESA components. To aid in that understanding, the following discussion draws an analogy to a more commonly understood architectural model—designing a house. It starts with a brief comparison of the house design model to the enterprise security system design model.

The House Design Model

- Community standards: the specific external and internal standards required by the housing community.
- Design requirements: the specific design criteria that are settled on after considering wants, needs, costs, etc. such as passive solar design with star wiring topology (LAN/telephony) and home entertainment system, plus fully handicapped-accessible downstairs with master bedroom and utilities.
- Building codes and engineering practices: the building standards and practices that support the design requirements and the architecture.
- Architectural plan: this is the resulting set of artist's renderings and blueprints that document what the house will look like from various perspectives. Also a necessary part of the plans are the detail drawings for the major components of the overall construction such as framing and the plumbing, electrical, and HVAC systems.
- Bill of materials: the detailed list of materials needed to build the house.
- Maintenance: the specific considerations for keeping the house up and its systems operational. Although not typically a significant part of the house design process, these specifications are relevant.

The Enterprise Security System Design Model

- Corporate standards: the specific corporate standards that affect the enterprise security system.
- Design requirements: the specific design criteria that are settled on after consideration of wants, needs, costs, etc One of these is already specified, the policy-driven security services. Other examples might include support for service-oriented application designs and role-based access control.
- Governance: the principles, policies, and implementing standards that support the design requirements and the specific architecture.
- Architectural plan: this is the resulting set of conceptual diagrams and blueprints that document what the resulting security system will look like from various perspectives. The plan includes the conceptual and detail drawings for major subsystems such as identity management, access control, and border protection services, as well as the required products, applications, platforms, etc.
- Security services: the itemization of the services and ultimately the individual applications and products needed.

- Operations: the considerations for day-to day-operation of the security services and supporting infrastructure.

Let’s take a look at each of these in detail and compare and contrast the components of the two models.

Community Standards vs. Corporate Standards

It’s important to keep in mind that both designs take place in a larger context that may impose constraints on the design—the house is part of a larger residential development or community, and the enterprise security system is part of a larger enterprise IT system.

In the house example, the community may impose standards to maintain a certain level of quality and appearance. It may, for example, restrict the use of certain types of siding and certain colors, and it may require a Jacuzzi and ceramic tile floor in the master bath and wood floors in certain rooms.

In the security example, corporate standards may be imposed to ensure that investments leverage existing technology or support infrastructure. They may, for example, require that all user-interfacing products support lightweight directory access protocol (LDAP) interoperability with their standard network operating system (NOS) or corporate directory to avoid proliferation of additional user registries and sign-on requirements.

Building Codes and Engineering Practices vs. Governance

In both models, development of the architectural plan must consider the constraints imposed by this component, based on experience and good judgment.

In the house example, building codes and engineering practices are constraints developed through years of experience to ensure a sound and safe dwelling. Considerations here include such things as structural integrity, a healthy environment, and fire safety. The finished architectural plan for each house may vary widely, but all must comply with these requirements.

In the security example, governance defines the principles, policies, standards, guidelines and procedures that constrain the design and operation of the security system. As with the house example, the governance elements are based on experience and good sense. Considerations include such things as simplicity, defense in depth, resilience, and common policy enforcement. As with the house example, security infrastructure implementations may vary widely, but all should comply with these requirements.

House Architecture vs. Security Technology Architecture

In both cases the architectural plan represents the blueprints for implementation. In the house example, the industry has built enough houses to clearly understand the various levels of detail and perspectives necessary for successful construction. Unfortunately, not many security infrastructures have been built using a comprehensive plan, so we are not nearly so clear on the levels of detail or perspectives needed.

One thing we do seem clear on is that any good plan starts with some high-level pictures and successively expands the detail in some organized fashion until the physical construction blueprints have been completed and construction can begin. In the computing industry these levels of detail are commonly termed the conceptual, logical, and physical architectures.

At the conceptual level, our design has artist's renderings of various views of the house. We see what the house looks like, possibly from various perspectives, but without any of the construction details or internal system components. In the security context, this should be a picture or pictures of the infrastructure as a whole, defining the key design concepts—hence, a conceptual architecture.

At the logical level, our house design has floor plans to specify the layout of each floor and show how the rooms are connected. There is still no detail of construction or the systems such as plumbing, heating, or framing. In the security design, this is where we see major services (such as identity management, access control, and border protection) decomposed into a set of related components and supporting services. For identity management, we see provisioning services, external and internal directories, policy administration systems, HR systems, identity mapping services, and more.

At the physical level, our house design has details for assembling the framing, electrical, plumbing, and HVAC components. In the security context, we see deployment of products and applications that make up the various functional components; we see computing platforms and connectivity.

Bill of Materials vs. Security Services

The security services are the security infrastructure bill of materials. These are the core functions we need to actually assemble a cohesive security infrastructure. To better understand this area, it's useful to look at the similarities between a house bill of materials and security services.

In both cases it is easy to start by itemizing a high-level bill of materials. We all know what kinds of material it takes to build a house. We need lumber, concrete, pipes, fixtures, ducting, fasteners, etc. We can easily make this list, but without the detailed plan we are not able to specify the quantities and types of each component. Similarly, we all know what security services are needed, but without the plan we cannot accurately list the specific products and platforms. The bill of materials is not an integral part of the plan, although it is a necessary part of the overall effort. The detailed bill of materials is *derived* from the plan. The list of security services at the detailed product level allows us to know what we need to build or buy to implement our plan.

Although natural, it is a mistake to think we can start with this bill of materials (list of security services) and somehow derive the plan (this is discussed a little more under the heading of "The Remodeling" below).

Maintenance vs. Operations

Once we have completed our house or security infrastructure, we need some processes and tools to maintain our work in a quality state. Furthermore, we probably need to take maintenance requirements into account in the design phase to facilitate our maintenance activities after completion.

In our house example, design elements related to maintenance might include selection of siding and flooring materials, installation of a built-in vacuum system, or placement of hose bibs to facilitate washing exterior components. Typical maintenance considerations after construction might be a daily cleaning plan, periodic painting and structural repair, regular heating and plumbing maintenance, and an occasional upgrade or addition.

In the security context, operations includes processes and tools for day-to-day vulnerability management, event management, and incident management, as well as other aspects of daily security administration and operation. These elements ensure continued effective and efficient functioning of the security environment.

The Remodeling

Most enterprises do not start with a green field in the security infrastructure space. We all have existing environments developed over the years, typically started with independent proprietary platforms, each with its own security silo. The advent of the Internet has been the primary driver for the deployment of a variety of products and solutions that attempt to integrate these disparate systems. For most of us the current state is a hodge-podge of environments and tools in various states of interoperability. The good news is:

- These point and reactive solutions have been built by smart people. Even if they did not use a comprehensive plan, these smart people typically made decisions and deployed solutions with a vision of what the plan should ultimately look like.
- There is increasing focus on the development of security standards to deliver interoperability among these disparate platforms. Many standards are in the early stages of development or adoption, so it is likely that interoperability will improve with time. At the same time, however, the interoperability challenge is increasingly complex.
- For the solutions we already have deployed, the marketplace is driving the vendors to continually enhance their interoperability, thus making our lives easier.

What this means as we work to articulate our new enterprise security infrastructure design is that we already have much of our bill of materials and we can probably use a substantial portion of our existing deployment.

So in the context of our analogy, we are possibly talking about house remodeling, not new construction. Somewhat contrary to what was stated earlier, the bill of materials will not be completely derived from the plan. There will now be some consideration of the existing construction (and inherent bill of materials) incorporated into our new design. However, in this case it is probably not wise to overemphasize the existing deployment when laying out the conceptual and upper-level aspects of the logical design. Consideration of the existing infrastructure will have more influence on the details of the logical design subcomponents and the physical design.

In our security context this remodeling probably means:

- Leveraging existing work to identify the security drivers and governance components. Care should be taken to be comprehensive at this point in the effort and not to assume that previous work is up to date in our rapidly changing environment.
- Assessing our existing environment and products as we work through the lower-level logical design and physical design. Much of what we have should be usable in our new comprehensive vision.
- Identifying gaps and areas for improvement in our existing infrastructure and then making plans for closing the gaps and implementing the improvements.

With the house analogy as background, let's move on to describe the ESA framework and templates, starting with security governance and then describing security technology architecture and security operations. Hopefully the house analogy has provided a basis for clearer understanding of some of the terms we use, and at appropriate points, we'll refer to the analogy again to clarify the discussion.

Security Governance

The focus now shifts to the security governance components and processes of NAC's overall framework, shown in the left center of Figure 5.

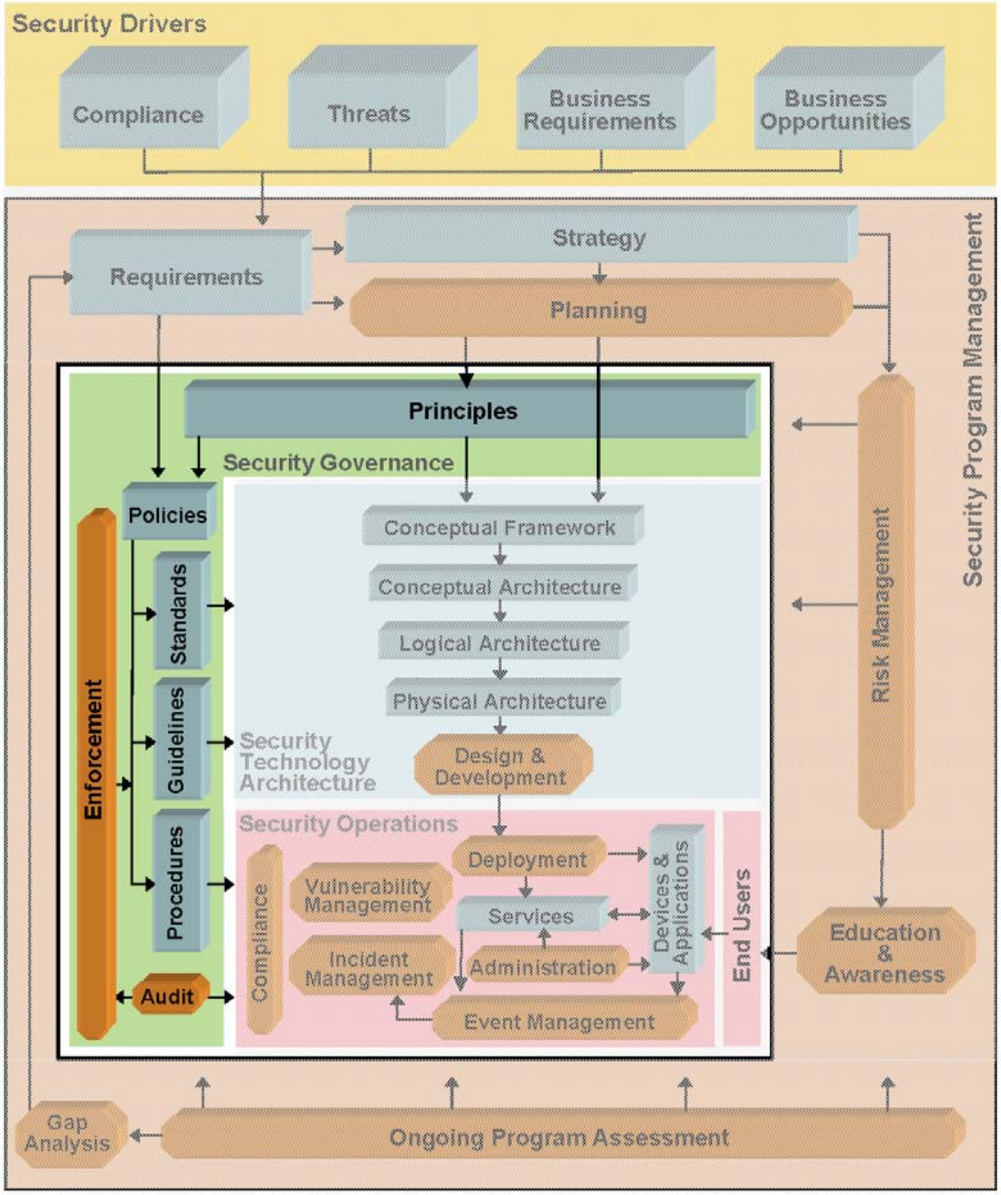

Figure 5. Security Governance Components and Processes

This section provides an overall security governance framework and template that member organizations can tailor to their needs. The governance components and processes were introduced earlier as follows:

- Principles: basic assumptions and beliefs providing overall security guidance.
- Policies: the security rules that apply in various control domains.

- Standards, guidelines, and procedures: the implementation of the policies through technical requirements, recommended practices, and instructions.
- Audit: the process of reviewing security activities for policy compliance.
- Enforcement: the processes for ensuring compliance with the policies.

The following sections provide an overview of the overall governance process and the policy framework, followed by descriptions of the individual components and processes identified above.

Governance Process Overview

For technicians, it is easy enough to find technical solutions to business problems. For example, there are various solutions for protecting a customer's identity, but how do technicians know their responsibility is to protect identity? How do they know that management has mandated this requirement? How do they know what the standards and guidelines are for implementing this requirement? How do they know what resources and services are available to implement a solution? Or more simply put, how do you, as a technician in an IT service organization, know what needs to be done to provide and maintain secure technical solutions that support the business mission and objectives of your organization?

NAC has identified this critical ESA component as governance. In NAC's vision of ESA, there is a strong linkage among governance, technology architecture, and operations. That linkage is provided via the *policy framework*, which is at the heart of the governance model, and the *policy-driven security architecture framework*, which is at the heart of the technology architecture and operations model. Before we describe the policy framework, it's useful to look at the overall governance process.

If we take a process-oriented view of defining a governance framework, the first step is to *identify* the guiding principles that your organization will follow in securing the information technology assets of the enterprise. These principles provide the highest level of guidance for the security governance process as well as for technology architecture and operations.

The second step is to *authorize* enforcement of the guiding principles through the creation of policies in various domains of management control. The control domains—such as organizational security, asset classification and control, personnel security, and access control—represent the highest-level identification of policy. The specific policies within each of these domains authorize a course of action.

The third step is to *implement* the authorized courses of action. The results are the technical standards, guidelines, and procedures that govern IT security for the organization.

The two additional governance concepts are *enforcement* and *ongoing assessment*. Typically, enforcement controls are built into the technical standards and procedures, but there are also requirements for separate enforcement processes triggered, for example, as a result of security-related events. Ongoing assessment is needed to respond to change as business models evolve, new technologies are developed, and new legislation is passed. An example of such a change occurred in the 1990s when business products and services were suddenly offered directly to the consumer through Web-based front ends to the traditional services. This change created the need to extend confidentiality principles to encompass the protection of personal data—the need for privacy protection is now taken for granted and is in many cases mandated by law. The effects of this sea change are still unfolding through the implementation of the Health Insurance Portability and Accountability Act (HIPAA) and other privacy-related legislation. Ongoing assessment is necessary to detect and respond to smaller changes as well and should be a built-in process for continuous improvement.

Governance Process Roles

Many different people are involved in identifying the guiding principles, authorizing them through policies, implementing and enforcing the policies, and continually assessing the effectiveness of the governance process. These people are not only involved in creating and maintaining the governance framework but may also have roles in technology architecture and operations.

Organizational managers are responsible for defining the organization’s principles by classifying the data used to drive the organization’s business needs based on legal, statutory, regulatory and contractual agreements. Organizational management is also responsible for managing risk.

Information systems management or the CIO is responsible for managing an organization’s technical systems that support the business services identified by organizational management through the creation and maintenance of policies.

The *security officer* is responsible for the security of a company's communications and other business systems. The security officer may also work with the CIO in planning for and managing disaster recovery. The security officer is likely to be involved in both the business (including people) and technical aspects of security, and is responsible for managing security incidents.

The *data security officer* assists with identifying and assessing risks associated with an organization’s data structure. This includes how data is accessed, stored, managed, and transferred.

Technical architects are responsible for building policy enforcement into the technical architecture.

Technicians (operations) apply the standards, guidelines, and procedures to their areas of responsibility.

Policy Framework Overview

At the heart of the governance model is the *policy framework*. As mentioned earlier, NAC’s vision of ESA includes a strong linkage among governance, technology architecture, and operations. At the governance level, the *policy framework* provides this linkage.

Figure 6 identifies NAC's generic policy framework. The basic framework concept is very simple—however, concept simplicity does not necessarily provide ease of definition and implementation.

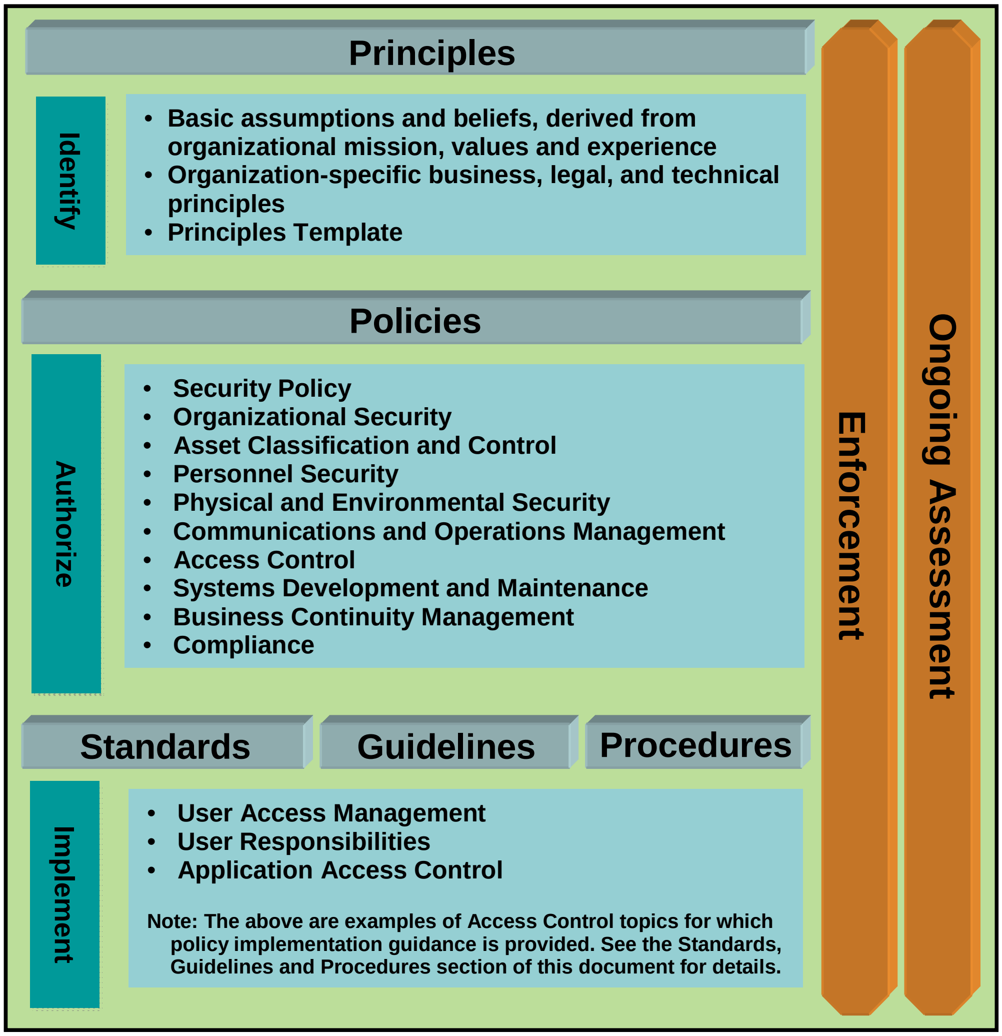
Principles
Identify
• Basic assumptions and beliefs, derived from organizational mission, values and experience
• Organization-specific business, legal, and technical principles
• Principles Template
Policies
Authorize
• Security Policy
• Organizational Security
• Asset Classification and Control
• Personnel Security
• Physical and Environmental Security
• Communications and Operations Management
• Access Control
• Systems Development and Maintenance
• Business Continuity Management
• Compliance
Standards
Guidelines
Procedures
Implement
• User Access Management
• User Responsibilities
• Application Access Control
Note: The above are examples of Access Control topics for which policy implementation guidance is provided. See the Standards, Guidelines and Procedures section of this document for details.
Enforcement
Ongoing Assessment

Figure 6. Generic Policy Framework

The following discussion explains the details of the framework.

- *Identify* the guiding principles for your organization:
  - Start with the fundamental objectives of IT security: availability, integrity, confidentiality, accountability, and assurance.
  - *Identify* basic assumptions and beliefs, derived from your organization's mission, values and experience.
  - *Identify* organization-specific business, legal, and technical principles.
  - *Tailor* the NAC principles template to your needs, based on these and any other organization-specific considerations.
- *Authorize* enforcement of your organization's guiding principles through an agreed-upon policy template:

    - Start with the *ISO/IEC 17799 policy template*, keeping in mind that these are high-level guidelines, not detailed technical guidelines.
    - Modify and extend the policy template based on your guiding principles and business needs.
    - Or purchase an *ISO/IEC 17799-compliant set of policies* and modify them as required to align with guiding principles and business needs.
- *Implement* the standards, guidelines, and procedures for your organization's technical environment, based on your policy template:
    - The Standards, Guidelines, and Procedures section includes examples of policy implementation guidance from the ISO/IEC 17799 document.
    - The Security Governance Resources and Tools section will identify additional sources of implementation guidance.
- *Enforce* compliance with policy. Enforcement is typically built into the technical standards and procedures, and it is supported by the NAC's policy-driven security architecture (see Security Technology Architecture). In addition, there are requirements for separate enforcement processes triggered as a result of security-related incidents or audits. *Audit* is called out in NAC's overall security program framework diagram because of its importance in supporting the requirement for accountability to the individual level.
- *Conduct ongoing assessments* for evaluating and responding to changes that may impact security policy—e.g., when business requirements change, new threats arise, new technologies are developed, and new legislation is passed.

ISO/IEC 17799:2000, *Code of Practice for Information Security Management*, is an international standard that is gaining traction in the enterprise security space. NAC selected it as an integral part of the policy framework based on industry recommendations received at the 2004 Spring Conference. It is believed to have broad applicability across the many organizational types represented in NAC. The British Standards Institution (BSI) has granted NAC permission to use the extracts identified in the policy framework description and reproduced in detail in the Policy Template and Standards, Guidelines, and Procedures sections that follow.

Principles

Technology governance principles are the basic assumptions, beliefs, theories, and values guiding the use and management of technology within an organization. All policies, standards, architectures, designs, operations, and other components of the technology process should align with these principles unless a governance body grants an exception.

Depending on the organization, governing principles may be established at one or more levels; this document focuses on the governing principles for enterprise security. *Identifying* an organization's guiding security principles is the first critical step in the governance process. These security principles constrain the definition of the other governance components, such as policies and standards, and they constrain the definition of the technology architecture and operations components. As NAC members adapt these principles to the needs of their particular organizations, they must ensure alignment with their higher-level corporate IT principles, which provide guidance on the use and deployment of all IT resources and assets across the enterprise.

Principles Template

The following NAC definition of security principles is based on input from several member organizations as well as NIST[5] and Microsoft[6]. This input was sorted into eight categories that represent the highest-level principles. Within each category are second-order and in some cases third-order principles.

Security by Design

Security should not be an afterthought or add-on. Security considerations should begin with the requirements phase of development and be treated as an integral part of the overall system design.

- Establish a sound security policy as the “foundation” for design.
- Build security into the life cycle.
    - Plan for system maintenance.
    - Ensure proper security in the shutdown or disposal of a system.
    - Commit to secure operations.
- Clearly delineate the physical and logical security boundaries governed by associated security policies.
- Protect technology assets through a comprehensive security program that includes appropriate security education, processes, and tools.
    - Invest in secure design.
    - Train developers in the techniques, processes, and tools needed to ensure secure software.
- Define the organizational roles and responsibilities required to implement security by design in your culture.

Managed Risk

Risk and security countermeasures should be balanced according to business objectives. Identify potential trade-offs between reducing risk and increasing cost, including negative impacts on other aspects of operational effectiveness, if any.

- Reduce risk to an acceptable level.
- Identify and prevent common errors and vulnerabilities.

[5] SP 800-27 Engineering Principles for Information Technology Security (A Baseline for Achieving Security), Revision A

[6] Security at Microsoft, Technical White Paper, Published: November 2003

- Assume that external systems are insecure.
- Ensure that the cost of security controls does not exceed the benefits (i.e., the tangible and intangible costs of the losses that could occur in the absence of the controls).

Usability and Manageability

Two aspects of usability must be considered—the end-user experience and the ease of administration and operation. Security should be user transparent and not cause users undue extra effort. Administration and configuration of security components should not be overly complex or obscure.

- Base security on open standards for portability and interoperability.
- Use common language in developing security requirements.
- Design security to allow for regular adoption of new technology, including a secure and logical technology upgrade process.
- Automate identity and access management activities.
- Strive for operational ease of use.

Defense in Depth

Greater security is obtained by layering defenses.

- Ensure that there is not just a single point of protection.
- Implement security through a combination of measures distributed physically and logically.
- Isolate public access systems from mission-critical resources (e.g., data, processes, etc.).
- Use common boundary mechanisms to separate computing systems and network infrastructures.

Simplicity

Complexity is the enemy of security. Systems should be as simple as possible while retaining functionality.

- Minimize the number of system elements to be trusted. Reduce the attack surface.
- Do not implement unnecessary security mechanisms.
- The number of security modules and services in the corporate systems environment should be minimized based on technical feasibility, cost, and security requirements.

Resilience

Design and operate IT systems so as to limit vulnerability and to be resilient in response. Automated recovery from attack or failure is desirable. The design should include the ability to restore operations in the event of a disaster, within a timeframe appropriate to business needs.

- Take appropriate measures to secure the information and communications business-critical infrastructure to enable business continuity in the event of disaster or attack.
  - Exercise contingency or disaster recovery procedures to ensure appropriate availability.
  - Ensure that security systems support restoration of data and recovery of function.
- Protect against all likely classes of attacks.
- Ensure that the system is, and will continue to be, resilient in the face of expected threats.
- Security components should fail closed—in other words, failure should result in denial of access rather than increased accessibility[7].
- Limit or contain vulnerabilities.
- Build in availability and redundancy.

Integrity

All components of the computing environment must provide for information integrity and confidentiality.

- Protect information while it is being processed, in transit, and in storage.
- Protect personally identifiable information and enforce other privacy requirements.
- Base decisions on data classification and fair use.
- Protect resources by using strong authentication.
- Formulate security measures to address multiple overlapping information domains.
- Authenticate users and processes to ensure appropriate access control decisions both within and across domains.
- Design and implement audit mechanisms to detect unauthorized use and to support incident investigations.
- Monitor and audit system access and use.
- Practice incident response.
- Limit access to systems and data to the least privilege required to perform a job function.

[7] A component failure should result in no access being granted, as opposed to a failure leaving the system open to accidental or intentional access

- Permit external access to enterprise technology assets only through methods that ensure enforcement of appropriate security measures.

Enforced Policy

Implement processes, procedures, and systems that promote enforcement of organizational security policies. Design component configuration procedures in accordance with security policy. Automate access control decisions based on corporate user identity information and access control policy statements.

- Implement policy-driven access control.
- Monitor identity confirmation.
- Use unique identities to ensure accountability.
- Leverage common enterprise identity and access management services.
- Distribute management of identity information.
- Utilize role-based and/or policy-based access control for authorization.
- Enforce secure configuration and hardening.

Policies

Policies define the authorizations and a program of actions adopted by an organization to govern the use of technology in specific areas of management control. Policies are a security governance tool used to enforce an organization's guiding principles. They are established and maintained through standards, guidelines, and procedures in accordance with related legal and business principles.

The development, use, and enforcement of policies as well as the level of policy detail may differ among organizations based on their business functions, cultures, and technology models. One organization may have a few policies that authorize the creation of many standards, guidelines, and procedures, while another organization may embed standards, guidelines, and procedures within its policies. In addition, policy development, enforcement, and maintenance strategies may differ.

Policy Development

Policy development history and current practice vary widely among organizations. Even applying a model such as the ISO/IEC 17799 code of practice can be difficult unless the organizational goals are first identified. The following are questions worth considering when creating a policy:

Analysis Questions

- What is being protected?
- Which principle or principles does the policy enforce?
- To whom does the policy apply, and are there limitations in the policy?
- Does the policy fit the organization's business needs and culture?

- Does the policy relate to the activities that actually take place within the organization?
- What deviations from the policy are acceptable?
- Does the policy state what must be done and what happens if the policy is not carried out?

Implementation Questions

- Who approved, authorized, and deployed the policy?
- When does the policy take effect?
- If the policy ends, when?
- Does all appropriate management properly support the policy?

Enforcement Questions

- Is the policy enforceable?
- Who is responsible for enforcing the policy?
- What are the ramifications of noncompliance?
- Who is responsible for monitoring and reporting policy violations?

Maintenance Questions

- Who is responsible for updating and maintaining the policy?
- How often should the policy be reviewed and updated?

Communication Questions

- How is the policy communicated?
- How are changes to the policy communicated?

Policy Template

The following table identifies the policy standards included in the ISO/IEC 17799 code of practice for information security management. The policies are described at three levels, starting with sections 3, 3.1, and 3.1.1. The first level is what we have referred to as the policy domain. Note that the first domain is security policy, which defines the requirement to develop and implement an information security policy. Altogether there are 10 policy domains:

- Security Policy
- Organizational Security
- Asset Classification and Control
- Personnel Security
- Physical and Environmental Security
- Communications and Operations Management
- Access Control

- Systems Development and Maintenance
- Business Continuity Management
- Compliance

Extracts from BS EN ISO 17799: 2000 are reproduced as part of Enterprise Security Architecture: A Framework and Template for Policy Driven Security document with the permission of BSI under license number 2004AT0134. Hard copies of other British Standards are available from BSI Customer Services, 389 Chiswick High Road, London W4 4AL, United Kingdom. [Tel: + 44 (0)20 8996 9001]. E-mail: cservices@bsi-global.com

Table 1: ISO/IEC 17799 Policy Template

Section	Section Title	Standard
3	**SECURITY POLICY**	
3.1	***Information Security Policy***	Management direction and support for information security must be clearly established.
3.1.1	Information Security Policy Document Development	Develop an Information Security Policy.
3.1.2	Review and Evaluation	Implement an Information Security Policy.
4	**ORGANIZATIONAL SECURITY**	
4.1	***Information Security Infrastructure***	A management framework must be established to initiate and control the implementation of information security within the organization.
4.1.1	Management Information Security Forum	Establish a corporate committee to oversee information security.
4.1.2	Information Security Coordination	Develop and implement an Information Security Organization mission statement.
4.1.3	Allocation of Information Security Responsibilities	Identify the roles and responsibilities of the Information Security Organization.
4.1.4	Authorization Process for Information Processing Facilities	Establish a management approval process to authorize new IT facilities from both a business and technical standpoint.
4.1.5	Specialist Information Security Advice	Charge the IS Organization with providing specialized information security advice.
4.1.6	Cooperation Between Organizations	Establish a liaison requirement with external information security personnel and organizations including industry and/or government security specialists; law enforcement authorities; IT service providers; telecommunications authorities.
4.1.7	Independent Review of Information Security	Identify that independent reviews of information security practices are conducted to ensure feasibility, effectiveness, and compliance with written policies.
4.2	***Security of Third Party Access***	The organizational IT facilities and information assets that control the access of non-organizational third parties must be kept secure.

4.2.1	Identification of Risks from Third Party Access	Implement a process to analyze third party connection risks. Implement specific security standards to combat third party connection risks.
4.2.2	Security Requirements in Third Party Contracts	Ensure that security requirements are included in formal third party contracts.
4.3	***Outsourcing***	The security of information should be maintained even when the responsibility for the processing has been outsourced to another organization.
4.3.1	Security Requirements in Outsourcing Contracts	Implement standards to address security requirements of the information owners in a contract between the owners and any outsource organization.
5	**ASSET CLASSIFICATION & CONTROL**	
5.1	***Accounting of Assets***	Appropriate accounting of organizational assets must be established.
5.1.1	Inventory of Assets	Establish an inventory of major assets associated with each information system.
5.2	***Information Classification***	Ensure that information assets receive an appropriate level of protection.
5.2.1	Classification Guidelines	Implement standards for security classification for the level of protection required for information assets.
5.2.2	Information Labeling and Handling	Implement standards to ensure the proper handling of information assets.
6	**PERSONNEL SECURITY**	
6.1	***Security in Job Definitions and Resourcing***	Security should be addressed at the recruitment stage, included in job descriptions and contracts, and monitored during an individual's employment.
6.1.1	Including Security in Job Descriptions	Ensure that security responsibilities are included in employee job descriptions.
6.1.2	Personnel Screening and Policy	Implement standards to ensure that employment applications are screened for jobs that require access to sensitive information.
6.1.3	Confidentiality Agreement	Implement non-disclosure agreements for all employees and third parties.
6.1.4	Terms and Conditions of Employment	Implement standards to ensure that employee terms and conditions of employment include the employee’s responsibility for information security, including duration after employment and consequences of failure to fulfill these terms.
6.2	***User Training***	Users should be trained in security procedures and the correct use of IT facilities.
6.2.1	Information Security Education and Training	Implement training standards to ensure that users are trained in information security policies and procedures, security requirements, business controls and correct use of IT facilities.

6.3	***Responding to Security Incidents and Malfunctions***	Incidents affecting security should be reported through management channels as quickly as possible.
6.3.1	Reporting of Security Incidents	Implement procedures and standards for formal reporting and incident response action to be taken on receipt of an incident report.
6.3.2	Reporting of Security Weaknesses	Implement standards and procedures to ensure that users are aware of the requirement to note and report all observed or suspected security weaknesses in or threats to systems or services.
6.3.3	Reporting of Software Malfunctions	Implement standards and user training to ensure that users note and report to the proper location any software that does not function correctly.
6.3.4	Learning from Incidents	Implement standards to ensure mechanisms are in place to monitor the types, volumes, and costs of incidents and malfunctions.
6.3.5	Disciplinary Process	Update corporate discipline policies to include dealing with employees who violate security policies and procedures.
7	**PHYSICAL & ENVIRONMENTAL SECURITY**	
7.1	***Secure Areas***	IT facilities supporting critical or sensitive business activities belong in secure areas.
7.1.1	Physical Security Perimeter	Implement standards to ensure that physical security protection exists, based on defined perimeters through strategically located barriers throughout the organization.
7.1.2	Physical Entry Controls	Implement entry procedures to secure areas to ensure only authorized personnel can gain access.
7.1.3	Securing Offices, Rooms, and Facilities	Implement procedures for physical security for data centers and computer rooms that are commensurate with threats.
7.1.4	Working in Secure Areas	Implement standards and procedures to control personnel or third parties working in the secure area.
7.1.5	Isolated Delivery and Loading Areas	Implement standards to ensure that the computer room/data center delivery and loading areas are isolated to prevent unauthorized access.
7.2	***Equipment Security***	Equipment must be physically protected from security threats and environmental hazards.
7.2.1	Equipment Sitting and Protection	Implement standards to ensure that equipment is located properly to reduce risks of environmental hazards and unauthorized access.
7.2.2	Power Supplies	Implement procedures for electronic equipment to protect it from power failures and other electrical anomalies.
7.2.3	Cabling Security	Implement standards to protect power and telecommunications cabling from interception or damage.

7.2.4	Equipment Maintenance	Implement procedures to establish and correctly maintain IT equipment to ensure its continued availability and integrity.
7.2.5	Security of Equipment Off-Premises	Implement standards and procedures to ensure that equipment used off-site, regardless of ownership, is provided the same degree of protection afforded on-site IT equipment.
7.2.6	Secure Disposal or Reuse of Equipment	Information can be compromised through the careless disposal or reuse of equipment.
7.3	***General Controls***	Information and information processing facilities should be protected from disclosure to, modification of, or theft by, unauthorized persons, and controls should be in place to minimize loss or damage.
7.3.1	Clear Desk and Clear Screen Policy	Implement a clear desk/clear screen policy for sensitive material to reduce risks of unauthorized access, loss, or damage outside normal working hours.
7.3.2	Removal of Property	Implement procedures to ensure that personnel are required to have documented management authorization to take equipment, data or software off-site.
8	**COMMUNICATIONS AND OPERATIONS MANAGEMENT**	
8.1	***Operational Procedures and Responsibilities***	Responsibilities and procedures must be established for the management and operation of all computers and networks.
8.1.1	Documented Operating Procedures	Implement operating procedures to clearly document all that all operational computer systems are being operated in a correct, secure manner.
8.1.2	Operational Change Control	Implement procedures for controlling changes to IT facilities and systems to ensure satisfactory control of all changes to equipment, software, or procedures.
8.1.3	Incident Management Procedures	Implement standards and procedures to identify incident management responsibilities and to ensure a quick, effective, orderly response to security incidents.
8.1.4	Segregation of Duties	Implement standards and user training to ensure that sensitive duties or areas of responsibility are kept separate to reduce opportunities for unauthorized modification or misuse of data or services.
8.1.5	Separation of Development and Operational Facilities	Implement procedures to segregate development and production facilities to reduce the risk of accidental changes or unauthorized access to production software and data.
8.1.6	External Facilities Management	The use of an external contractor to manage information processing facilities may introduce potential security exposures such as the possibility of compromise, damage, or loss of data at the contractor's site.

8.2	***System Planning and Acceptance***	Advance planning and preparation can ensure the availability of adequate capacity and resources.
8.2.1	Capacity Planning	Implement standards to ensure that capacity requirements are monitored, and future requirements projected, to reduce the risk of system overload.
8.2.2	System Acceptance	Implement procedures to establish acceptance criteria for new systems, and to ensure that adequate tests have been performed prior to acceptance.
8.3	***Protection from Malicious Software***	Applying precautions to prevent and detect the introduction of malicious software can safeguard the integrity of software and data.
8.3.1	Controls Against Malicious Software	Implement standards and user training to ensure that virus detection and prevention measures are adequate.
8.4	***Housekeeping***	Routine procedures should be established for making back-up copies of data, logging events and faults, and where appropriate, monitoring the equipment environment.
8.4.1	Information Back-up	Establish procedures for making regular back-up copies of essential business data and software to ensure that it can be recovered following a computer disaster or media failure.
8.4.2	Operator Logs	Implement standards and procedures that computer operators are required to maintain a log of all work performed.
8.4.3	Fault Logging	Implement procedures for logging faults reported by users regarding problems with computer or communications systems.
8.5	***Network Management***	The security of computer networks that may span organizational boundaries must be managed to safeguard information and to protect the supporting infrastructure.
8.5.1	Network Controls	Implement appropriate standards to ensure the security of data in networks and the protection of connected services from unauthorized access.
8.6	***Media Handling and Security***	Computer media should be controlled and physically protected to prevent damage to assets and interruptions to business activities.
8.6.1	Management of Removable Computer Media	Implement procedures for the management of removable computer media such as tapes, disks, cassettes, and printed reports.
8.6.2	Disposal of Media	Implement standards and procedures to ensure that computer media is disposed of securely and safely when no longer required.
8.6.3	Information Handling Procedures	Implement procedures for handling sensitive data to protect such data from unauthorized disclosure or misuse.
8.6.4	Security of System Documentation	Implement standards to protect system documentation from unauthorized access.

8.7	***Exchanges of Information and Software***	Exchanges of data and software between organizations should be controlled to prevent loss, modification, or misuse of data.
8.7.1	Information and Software Exchange Agreements	Implement procedures to establish formal agreements exist, including software escrow agreements when appropriate, for exchanging data and software (whether electronically or manually) between organizations.
8.7.2	Security of Media in Transit	Implement standards to safeguard computer media being transported between sites to minimize its vulnerability to unauthorized access, misuse, or corruption during transportation.
8.7.3	Electronic Commerce Security	Implement standards to protect electronic commerce (electronic data interchange, electronic mail, and on-line transactions across a public network such as the Internet) against unauthorized interception or modification.
8.7.4	Security of Electronic Mail	Implement standards and user training to reduce the business and security risks associated with electronic mail to include interception, modification, and errors.
8.7.5	Security of Electronic Office Systems	Implement a risk analysis process and resultant standards to control business and security risks associated with electronic office systems.
8.7.6	Publicly Available Systems	Implement a formal policy to establish an authorization process for information that is to be made publicly available.
8.7.7	Other Forms of Information Exchange	Implement procedures and standards to protect the exchange of information through the use of voice, facsimile, and video communications facilities.
9	**ACCESS CONTROL**	
9.1	***Business Requirement for System Access***	Policies for information dissemination and entitlement should control access to computer services and data on the basis of business requirements.
9.1.1	Access Control Policy	Implement a risk analysis process to gather business requirements to document access control levels.
9.2	***User Access Management***	Formal procedures are needed to control allocation of access rights to IT services.
9.2.1	User Registration	Implement procedures for user registration and deregistration access to all multiuse IT services.
9.2.2	Privilege Management	Implement standards to protect against the use of any feature or facility of a multi-user IT system that enables a user to override system or application controls.
9.2.3	User Password Management	Implement standards to address password management.
9.2.4	Review of User Access Rights	Implement procedures to conduct periodic reviews of users' access rights.

9.3	***User Responsibilities***	Users should be made aware of their responsibilities for maintaining effective access controls, particularly regarding the use of passwords and security of user equipment.
9.3.1	Password Use	Implement user training to ensure that users have been taught good security practices in the selection and use of passwords.
9.3.2	Unattended User Equipment	Implement policies and procedures to ensure that all users and contractors are made aware of the security requirements and procedures for protecting unattended equipment. Implement standards to log that all users and contractors have been made aware of their responsibilities for implementing such protection.
9.4	***Network Access Control***	Connections to network services should be controlled to ensure that connected users or computer services do not compromise the security of any other networked services.
9.4.1	Policy on Use of Network Services	Implement procedures to ensure that network and computer services that can be accessed by an individual user or from a particular terminal are consistent with business access control policy.
9.4.2	Enforced Path	Implement standards that restrict the route between a user terminal and the computer services that its user is authorized to access.
9.4.3	User Authentication for External Connections	Implement standards to ensure that connections by remote users via public or non-organization networks are authenticated to prevent unauthorized access to business applications.
9.4.4	Node Authentication	Implement standards to ensure that connections by remote computer systems are authenticated to prevent unauthorized access to a business application.
9.4.5	Remote Diagnostic Port Protection	Implement procedures to control access to diagnostic ports designed for remote use by maintenance engineers.
9.4.6	Segregation in Networks	Implement standards to have large networks divided into separate domains to mitigate the risk of unauthorized access to existing computer systems that use the network.
9.4.7	Network Connection Control	Implement standards to restrict the connection capability of users, in support of access policy requirements of business applications that extend across organizational boundaries.
9.4.8	Network Routing Control	Implement standards that identify routing controls over shared networks across organizational boundaries to ensure those computer connections and information flows conform to the access policy of business units.

9.4.9	Security in Network Services	Implement standards to capture network providers security attributes of all services used, and use this information to establish the security controls to protect the confidentiality, integrity, and availability of business applications.
9.5	***Operating System Access Control***	Access to computers should be strictly limited through the use of: automatic terminal identification; terminal logon procedures; user IDs; password management; a duress alarm; terminal time out; and limited connection time.
9.5.1	Automatic Terminal Identification	Implement standards for automatic terminal identification to authenticate connections to specific locations.
9.5.2	Terminal Logon Procedures	Implement procedures for logging into a computer system to minimize the opportunity for unauthorized access.
9.5.3	User Identification and Authentication	Establish standards to ensure all users have a unique identifier (user ID) for their personal and sole use, to ensure that their activities can be traced to them.
9.5.4	Password Management System	Implement standards to ensure an effective password management system is employed to authenticate users.
9.5.5	Use of System Utilities	Implement standards to restrict access to system utility programs that could be used to override system and application controls.
9.5.6	Duress Alarm to Safeguard Users	Conduct a risk analysis to determine if a duress alarm needs to be provided for users who might be the target of coercion. Implement standards to define responsibilities for responding to duress alarms.
9.5.7	Terminal Time-Out	Implement standards to ensure that terminals in high-risk locations are set to time out when inactive to prevent access by unauthorized persons.
9.5.8	Limitation of Connection Time	Implement standards to identify the period during which terminals may be connected to sensitive application systems.
9.6	***Application Access Control***	Logical access controls should be enacted to protect application systems and data from unauthorized access.
9.6.1	Information Access Restriction	Implement procedures to restrict access to applications system data and functions in accordance with defined access policy and based on individual requirements.
9.6.2	Sensitive System Isolation	Implement standards to isolate sensitive application systems processing environments.
9.7	***Monitoring System Access and Use***	Systems should be monitored to ensure conformity with access policy and standards, to detect unauthorized activities, and to determine the effectiveness of security measures adopted.

9.7.1	Event Logging	Implement standards to have audit trails record exceptions and other security-relevant information and ensure that they are maintained to assist in future investigations and in access control monitoring.
9.7.2	Monitoring System Use	Implement procedures for monitoring system use to ensure that users are only performing processes that have been explicitly authorized.
9.7.3	Clock Synchronization	Implement standards to ensure computer or communications device clocks are correct and in synchronization.
9.8	***Mobile System Access and Use***	When using mobile computing and telecommuting, the organization should examine the risks and apply appropriate protection to the equipment or site.
9.8.1	Mobile Computing	Implement a formal policy and supporting standards that address the risks of working with mobile computing facilities, including requirements for physical protection, access controls, cryptographic techniques, back up, and virus protection.
9.8.2	Teleworking	Implement policies and procedures to control telecommuting, to include existing facilities, the proposed telecommuting environment, communications security requirements, and the threat of unauthorized access to equipment or the network.
10	**SYSTEMS DEVELOPMENT & MAINTENANCE**	
10.1	**Security Requirements of Systems**	To ensure that security is built into IT systems, security requirements should be identified, justified, agreed to, and documented as part of the requirements definition stage of all IT system development projects.
10.1.1	Security Requirements Analysis and Specification	Implement standards to ensure that analysis of security requirements is part of the requirement analysis stage of each development project.
10.2	**Security in Application Systems**	Security controls that conform to commonly accepted industry standards of good security practice should be designed into applications systems to prevent loss, modification, or misuse of user data.
10.2.1	Input Data Validation	Implement standards to ensure that data that is input into applications systems is validated to ensure that it is correct and appropriate.
10.2.2	Control of Internal Processing	Implement standards to ensure that validation checks are incorporated into systems to detect corruption caused by processing errors or through deliberate acts.
10.2.3	Message Authentication	Implement standards to ensure that message authentication is considered for applications that involve the transmission of sensitive data.

10.2.4	Output Data Validation	Implement standards to ensure that data that is output from applications systems is validated to ensure that it is correct and appropriate.
10.3	***Cryptographic Controls***	To protect the confidentiality, authenticity, or integrity of information, cryptographic systems and techniques should be used for complete protection of information that is considered at risk.
10.3.1	Policy on the Use of Cryptographic Controls	Implement policies and standards on the use of cryptographic controls, including management of encryption keys, and effective implementation.
10.3.2	Encryption	Implement standards to ensure that data encryption is used to protect highly sensitive data during transmission or in storage.
10.3.3	Digital Signatures	Implement standards for the use of digital signatures to protect the authenticity and integrity of electronic documents.
10.3.4	Non-repudiation Services	Implement standards for non-repudiation services where disputes might arise based on the use of encryption or digital signatures.
10.3.5	Key Management	Implement standards for use of cryptographic techniques, including secret key techniques and public key techniques.
10.4	***Security of System Files***	To ensure that IT projects and support activities are conducted in a secure manner, the responsibility for controlling access to application system files should be assigned to and carried out by the owning user function or development group.
10.4.1	Control of Operational Software	Implement standards to ensure that strict control is exercised over the implementation of software on operational systems.
10.4.2	Protection of System Test Data	Implement standards to ensure that all application system test data is protected and controlled.
10.4.3	Access Control to Program Source Library	Implement standards and procedures to restrict access to program source libraries to reduce the potential for corruption of computer programs.
10.5	***Security in Development and Support Processes***	Project and support environments must be strictly controlled to maintain the security of application system software and data.
10.5.1	Change Control Procedures	Implement standards and procedures for formal change control.
10.5.2	Technical Review of Operating System Changes	Implement procedures to review application systems when changes to the operating systems occur.
10.5.3	Restrictions on Changes to Software Packages	Implement standards to restrict modifications to vendor-supplied software.

10.5.4	Covert Channels and Trojan Code	Implement standards and procedures to avoid covert channels or Trojan codes These standards and procedures should address at a minimum that the organization: buy programs only from a reputable source; buy programs in source code that is verifiable; use only evaluated products; inspect all source code before operational use; control access to, and modification of, installed code; use trusted staff to work on key systems
10.5.5	Outsourced Software Development	Implement standards to address when software development is outsourced, that controls are in place to address the ownership of intellectual property throughout the project life cycle.
11	**BUSINESS CONTINUITY MANAGEMENT**	
11.1	***Aspects of Business Continuity Management***	Business continuity plans (BCPs) should be available to counteract interruptions to business activities.
11.1.1	Business Continuity Management Process	Implement procedures for the development and maintenance of business continuity plans across the organization.
11.1.2	Business Continuity and Impact Analysis	Implement standards for corporate business continuity, and ensure that management endorses the plan.
11.1.3	Writing and Implementing Continuity Plans	Implement standards and procedures for business continuity planning to encompass the identification of all responsibilities and emergency procedures.
11.1.4	Business Continuity Planning Framework	Implement a single business continuity plan framework maintained to ensure that all levels of the plan are consistent.
11.1.5	Testing, Maintaining, and Re-assessing Business Continuity Plans	Implement standards to ensure regular testing of the BCPs.
12	**COMPLIANCE**	
12.1	***Compliance with Legal Requirements***	All relevant requirements for each IT system should be identified and documented.
12.1.1	Identification of Applicable Legislation	Implement standards to ensure that all relevant statutory, regulatory, and contractual requirements are specifically defined and documented for each information system.
12.1.2	Intellectual Property Rights (IPR)	Implement standards to ensure there is compliance with legal restrictions on the use of copyright material ensuring that only software developed by the organization, or licensed or provided by the developer to the organization, is used.
12.1.3	Safeguarding of Organizational Records	Implement policies and standards to ensure that important organizational records are securely maintained to meet statutory requirements, as well as to support essential business activities.

12.1.4	Data Protection and Privacy of Personal Information	Implement standards to ensure that applications that process personal data on individuals comply with applicable data protection legislation.
12.1.5	Prevention of Misuse of Information Processing Facilities	Implement policies to ensure that IT facilities are used only for business purposes.
12.1.6	Regulation of Cryptographic Controls	Implement standards and procedures to ensure that legal advice is sought on the organization’s compliance with national and international laws on cryptographic controls.
12.1.7	Collection of Evidence	Implement standards and procedures to ensure that when conducting an investigation the rules for evidence are followed for admissibility, quality, and completeness.
12.2	***Reviews of Security Policy and Technical Compliance***	To ensure compliance of IT systems with organizational security policies and standards, compliance reviews should be conducted regularly.
12.2.1	Compliance with Security Policy	Implement standards to ensure that all areas within the organization are considered for regular review to ensure compliance with security policies and standards.
12.2.2	Technical Compliance Checking	Implement standards to ensure that IT facilities are regularly checked for compliance with security implementation standards.
12.3	***System Audit Considerations***	There should be controls over operational systems and audit tools during system audits to minimize interference to and from the system audit process, and to protect the integrity and prevent the misuse of audit tools.
12.3.1	System Audit Controls	Implement standards to ensure audits and activities involving checks on operational systems are carefully planned and arranged.
12.3.2	Protection of System Audit Tools	Implement standards and procedures to restrict access to system audit tools.

Standards, Guidelines, and Procedures

Policies are *implemented* through technical standards, guidelines, and procedures, which NAC distinguishes as follows:

- Standards are mandatory directives.
- Guidelines are recommended best practices.
- Procedures describe how to achieve the standard or guideline. Usually they are incorporated within the standards or guidelines. In some cases, separate procedures may be needed, for example to establish a process whereby independent business units comply with corporate policies or standards.

This section does not attempt to provide a complete template for standards, guidelines, and procedures that implement the ISO/IEC 17799 policies. Such a template could include hundreds of specific standards covering a broad range of infrastructure topics and platforms. *Instead, this section offers a few examples of* *implementation guidance from the Access Control section of the ISO/IEC 17799 code of practice, specifically 9.2 User Access Management, 9.3 User Responsibilities, and 9.6 Application Access Control.*

9.2 User Access Management

Objective: To prevent unauthorized access to information systems.

Formal procedures should be in place to control the allocation of access rights to information systems and services.

The procedures should cover all stages in the life cycle of user access, from the initial registration of new users to the final de-registration of users who no longer require access to information systems and services. Special attention should be given, where appropriate, to the need to control the allocation of privileged access rights, which allow users to override system controls.

9.2.1 User registration

There should be a formal user registration and de-registration procedure for granting access to all multi-user information systems and services.

Access to multi-user information services should be controlled through a formal user registration process, which should include:

a) Using unique user IDs so that users can be linked to and made responsible for their actions. The use of group IDs should only be permitted where they are suitable for the work carried out.
b) Checking that the user has authorization from the system owner for the use of the information system or service. Separate approval for access rights from management may also be appropriate.
c) Checking that the level of access granted is appropriate to the business purpose (see 9.1) and is consistent with organizational security policy, e.g. it does not compromise segregation of duties (see 8.1.4).
d) Giving users a written statement of their access rights.
e) Requiring users to sign statements indicating that they understand the conditions of access.
f) Ensuring service providers do not provide access until authorizations procedures have been completed.
g) Maintaining a formal record of all persons registered to use the service.
h) Immediately removing access rights of users who have changed jobs or left the organization.
i) Periodically checking for, and removing, redundant user IDs and accounts.
j) Ensuring that redundant user IDs are not issued to other users.

Consideration should be given to including clauses in staff contracts and service contracts that specify sanctions if unauthorized access is attempted by staff or service agents (see also 6.1.4 and 6.3.5).

9.2.2 Privilege management

The allocation and use of privileges (any feature or facility of a multi-user information system that enables the user to override system or application controls) should be restricted and controlled. Inappropriate use of system privileges is often found to be a major contributory factor to the failure of systems that have been breached.

Multi-user systems that require protection against unauthorized access should have the allocation or privileges controlled through a formal authorization process. The following steps should be considered.

a) The privileges associated with each system product, e.g. operating system, database management system and each application, and the categories of staff to which they need to be allocated should be identified.
b) Privileges should be allocated to individuals on a need-to-use basis and on an event-by-event basis, i.e. the minimum requirement for their functional role only when needed.
c) An authorization process and a record of all privileges allocated should be maintained. Privileges should not be granted until the authorization process is complete.
d) The development and use of system routines should be promoted to avoid the need to grant privileges to users.
e) Privileges should be assigned to a different user identity from those used for normal business use.

9.2.3 User password management

Passwords[8] are a common means of validating a user's identity to access an information system or service. The allocation of passwords should be controlled through a formal management process, the approach of which should:

a) Require users to sign a statement to keep personal passwords confidential and work group passwords solely within the members of the group (this could be included in the terms and conditions of employment, see 6.1.4);
b) Ensure, where users are required to maintain their own passwords, that they are provided initially with a secure temporary password which they are forced to change immediately. Temporary passwords provided when users forget their password should only be provided following the positive identification of the user;
c) Require temporary passwords to be given in a secure manner. The use of third parties or unprotected (clear text) electronic mail messages should be avoided. Users should acknowledge receipt of passwords.

Passwords should never be stored on computer system in unprotected form (see Other technologies for user identification and authentication, such as biometrics, e.g. finger-print verification, signature verification and use of hardware tokens, e.g. chip-cards, are available, and should be considered if appropriate).

9.2.4 Review of user access rights

To maintain effective control over access to data and information services, management should conduct a formal process at regular intervals to review users' access rights so that:

[8] NAC members recognize that passwords are not particularly strong authenticators. The industry trend is towards more secure authentication technologies, such as smart cards or biometrics

a) Users' access rights are reviewed at regular intervals (a period of 6 months is recommended) and after any changes (see 9.2.1);
b) Authorizations for special privileged access rights (see 9.2.2) should be reviewed at more frequent intervals; a period of 3 months is recommended;
c) Privilege allocations are checked at regular intervals to ensure that unauthorized privileges have not been obtained.

9.3 User Responsibilities

Objective: To prevent unauthorized user access.

The cooperation of authorized users is essential for effective security. Users should be made aware of their responsibilities for maintaining effective access controls, particularly regarding the use of passwords and the security of user equipment.

9.3.1 Password use

Users should follow good security practices in the selection and use of passwords.

Passwords provide a means of validating a user's identity and thus to establish access rights to information processing facilities or services. All users should be advised to:

a) Keep passwords confidential;
b) Avoid keeping a paper record of passwords, unless this can be stored securely;
c) Change passwords whenever there is any indication of possible system or password compromise;
d) Select quality passwords with a minimum length of six characters which are
    1. Easy to remember;
    2. Not based on anything somebody else could easily guess or obtain using person related information, e.g. names, telephone numbers, and dates of birth, etc.;
    3. Free of consecutive identical characters or all-numeric or all-alphabetical groups.
e) Change passwords at regular intervals or based on the number of accesses (passwords for privileged accounts should be changed more frequently than normal passwords), and avoid re-using or cycling old passwords;
f) Change temporary passwords at the first log-on;
g) Do not include passwords in any automated log-on process, e.g. stored in a macro or function key;
h) Do not share individual user passwords.

If users need to access multiple services or platforms and are required to maintain multiple passwords, they should be advised that they may use a single, quality password [see d) above] for all services that provide a reasonable level of protection for stored password.

9.3.2 Unattended user equipment

Users should ensure that unattended equipment has appropriate protection. Equipment installed in user areas, e.g., workstations or file servers, may require specific protection from unauthorized access when left unattended for an extended period. All users and contractors should be made aware of the security requirements and procedures for protecting unattended equipment, as well as their responsibilities for implementing such protection. Users should be advised to:

a) Terminate active sessions when finished, unless they can be secured by an appropriate locking mechanism, e.g., a password protected screen saver;
b) Log-off mainframe computers when the session is finished (i.e., not just switch off the PC or terminal);
c) Secure PCs or terminals from unauthorized use by a key lock or an equivalent control, e.g., password access, when not in use.

9.6 Application Access Control

Objective: To prevent unauthorized access to information held in information systems. Security facilities should be used to restrict access within application systems. Logical access to software and information should be restricted to authorized users. Application systems should:

a) Control user access to information and application system functions, in accordance with a defined business access control policy;
b) Provide protection from unauthorized access for any utility and operating system software that is capable of overriding system or application controls;
c) Not compromise the security of other systems with which information resources are shared;
d) Be able to provide access to information to the owner only, other nominated authorized individuals, or defined groups of users.

9.6.1 Information access restriction

Users of application systems, including support staff, should be provided with access to information and application system functions in accordance with a defined access control policy, based on individual business application requirements and consistent with organizational information access policy (see 9.1). Application of the following controls should be considered in order to support access restriction requirements:

a) Providing menus to control access to application system functions;
b) Restricting users' knowledge of information or application system functions which they are not authorized to access, with appropriate editing of user documentation;
c) Controlling the access rights of users, e.g., read, write, delete, and execute;
d) Ensuring that outputs from application systems handling sensitive information contain only the information that are relevant to the use of the output and are sent only to authorized terminals and locations, including periodic review of such outputs to ensure that redundant information is removed.

9.6.2 Sensitive system isolation

Sensitive systems might require a dedicated (isolated) computing environment. Some application systems are sufficiently sensitive to potential loss that they require special handling. The sensitivity may indicate that the application system should run on a dedicated computer, should only share resources with trusted application systems, or have no limitations. The following considerations apply.

a) The sensitivity of an application system should be explicitly identified and documented by the application owner (see 4.1.3).

b) When a sensitive application is run in a shared environment, the application systems with which it will share resources should be identified and agreed with the owner of the sensitive application.

This is the end of the extracts from BS EN ISO 17799: 2000.

Enforcement

Enforcement is the overall process of ensuring compliance with policy. It is accomplished through a combination of technical controls, process and procedure controls, and management controls. Many of these controls are built into the implementing technical standards and procedures. Management controls provide for discretionary invocation of enforcement processes (such as disciplinary actions) as a result of security events or incidents. Management enforcement depends upon maintaining accountability for user actions, which is performed primarily by the audit and non-repudiation services.

Ongoing Assessment

Ongoing assessment is the process of evaluating and responding to changes that may impact any aspect of the governance process and policy framework.

Changes in business, legal, and technical principles need to be reviewed periodically in order to determine whether additions or modifications to security policy may be implied or even mandated.

Policies need to be reviewed to ensure that they are effectively protecting IT assets as intended. For example, if a security incident indicates that an unauthorized person was able to access data from an unattended workstation, then the policy that restricts inappropriate access needs to be reviewed for enforcement practices. The applicable standard also needs to be reviewed.

Standards, guidelines, and procedures also need to be reviewed on an ongoing basis as new employees are hired, new systems or services are implemented, or current systems are upgraded.

An effective and efficient ongoing assessment process requires supporting tools and metrics. The key is to collect and measure data that identifies the strengths and weaknesses of the security architecture as implemented. For example, it will be easier to demonstrate to management that a stronger anti-virus packages is required if you have historical metrics showing the impact of virus attacks on your organization.

Governance Example

This section provides an example from the fictitious XYZ Company to illustrate the relationship between security principles and policies and the implementing standards. In summary, the key relationships are:

- Policies define the rules for a particular domain—in this case, the domain is authentication passwords.

- Policy rule definitions must be consistent with guiding principles—in this case there are two guiding principles, integrity and usability.
- A policy may be implemented by multiple standards covering different aspects of the policy—in this example, only one of the standards is shown.

The names of the principles, policy, and standard for this example are shown in bold.

Integrity Principle

All components of the computing environment must provide for information integrity and confidentiality. Resources must be protected using strong authentication.

Usability and Manageability Principle

Two aspects of usability must be considered—the end-user experience and the ease of administration and operation. *Ideally, security should be user transparent and not cause users undue extra effort.* Administration and configuration of security components should not be overly complex or obscure.

Authentication Policy Example

Requirement:

Authentication mechanisms must be protected commensurate with the value of the information or business process they support, and they must be resistant to common methods of compromise.

Overview

Authentication mechanisms substantiate a claim of identity through the use of authentication data. All components of authentication systems need to be protected from unauthorized disclosure and misuse to preserve the integrity of the authentication. Additionally, the data components of authentication systems need to be protected commensurate with the sensitivity of the assets they help protect. The following components and processes accomplish authentication and protect authentication mechanisms:

- Selection of authentication data.
- Collection of authentication and verification data.
- Association of collected authentication and verification data with an identity.
- Protection of user's authenticator.
- Transport of the authenticator.
- Protection of verification data in storage.
- Validation of authenticator.
- Access permission or denial based on results of the authentication.

When authentication data is presented to the authentication mechanism, it is called an authenticator. Authenticators are generated in the following ways (called factors):

- Something you have (e.g., a token card).

- Something you know (e.g., a password).
- Something you are (e.g., fingerprints).

The three XYZ Company standard authentication types are designated as normal, supplemented, and strong.

Normal authentication is the weakest type allowed on the XYZ Company internal network. It includes clear text passwords and digital certificates where the private key protection can't be guaranteed.

Supplemented authentication is normal authentication that is implemented in combination with additional controls. Supplemented authentication is resistant to common methods of compromise and needs to be used instead of normal authentication when additional risk is present. For example, additional risk is present when an entity connects to an XYZ Company asset from outside the XYZ Company internal network. A standard way to implement supplemented authentication is using normal authentication within an approved encrypted channel such as a secure sockets layer (SSL).

Strong authentication is authentication that provides a high degree of accountability and assurance of identity on its own. To be considered strong, authentication needs to incorporate two factors. Standard ways to implement strong authentication include:

- Smart card with a personal identification number (PIN).
- Two-factor token card (e.g., a SecureID device such as a token that provides an ever-changing password, in combination with a PIN).
- X.509 certificate, where the private key is protected by a PIN that complies with the password selection standard.

Verification data is the information a system compares with the user's authenticator to validate the user's identity. Usually, verification data is not stored in the same format as the authenticator. Typically, a secure mathematical operation (encryption, one-way hash) is performed to derive verification data from the authenticator.

Password Quality Enforcement Standard Example

Systems must be configured so that static passwords cannot be reused for an account for eight reset cycles. When systems do not support eight cycles, the maximum number of cycles permitted by the system must be used.

Periodic checking for weak passwords should be performed. Password checking by any organization or individual must be authorized by Enterprise Computing Security.

- When a weak password is discovered, the user should be notified to change the password immediately. If the user is unavailable or unable to comply, the account should be disabled.
- System administrators desiring to discover weak passwords must confine their activities to systems under their direct responsibility.
- Any password scanning software or resulting data must be protected, and visibility must be limited to persons with a need to know.

Systems should be configured to enforce password complexity, when such capability is provided by the infrastructure.

Systems configured to enforce password complexity should allow passwords that can be used on multiple systems. Such passwords have the following characteristics:

- They contain eight characters.
- The first character is alphabetic.
- At least one numeric character is in the second through seventh character position.
- Last character is non-numeric.
- When the system is case sensitive, the password includes at least one uppercase and one lowercase character.

Example Comments

This example illustrates an NAC member's implementation of authentication policy through a password quality enforcement standard. Not shown are 13 other standards that implement various aspects of the NAC member's authentication policy.

The password quality enforcement standard stipulates periodic checking for weak passwords and mandates their replacement. The standard also takes into account the principle that *security should be user transparent and not cause users undue extra effort* by allowing for passwords that can be used on multiple systems.

Also note that in this example enforcement is built into the standard—if the user is unavailable or is unable to comply, the user account is disabled. Other examples may not be so simple or clear cut and may involve a separate enforcement process that invokes disciplinary actions.

Security Technology Architecture

The focus now shifts to the security technology architecture components and processes of NAC's overall framework, shown in the center of Figure 7.

The purpose of this section is to provide an overall security technology architecture framework and template that member organizations can tailor to their needs. The overall framework is described at four levels of abstraction: conceptual framework, conceptual architecture, logical architecture, and physical architecture. These were introduced briefly during the earlier discussion of the security program framework, as follows:

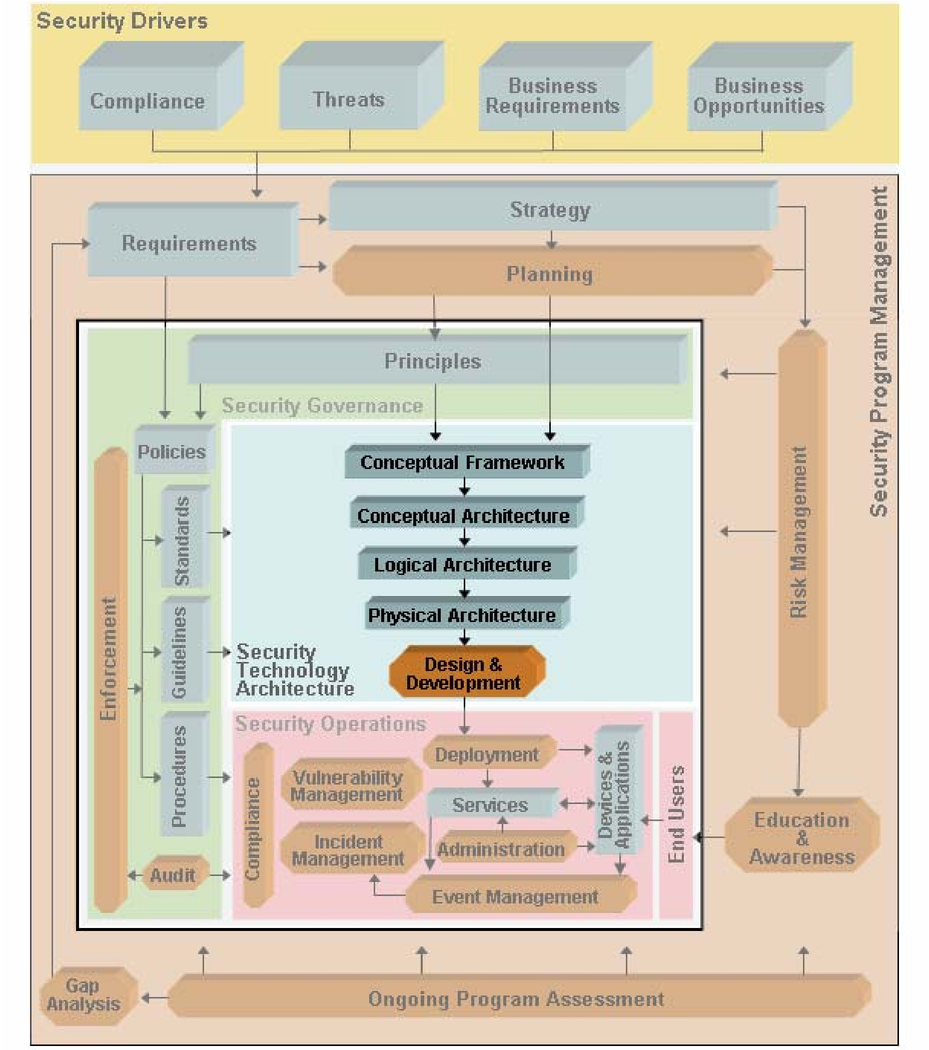

Figure 7. Security Technology Architecture Components and Processes

- Conceptual framework: generic technical framework for policy-based security services.
- Conceptual architecture: conceptual structure for management of decision making and policy enforcement across a broad set of security services.

- Logical architecture: structure and relationships of the key components and services defined within the constraints of the conceptual architecture.
- Physical Architecture: identifies structure of specific products, showing their placement and the connectivity relationships required to deliver the necessary functionality, performance, and reliability within the constraints of the logical architecture.

These definitions are tied closely to NAC's vision of policy-driven security, with a strong linkage among governance, technology architecture, and operations. This is the fundamental concept underlying NAC's definition of ESA and forward-looking enterprise security system implementations. The following sections explain the concept, starting with the conceptual framework.

Conceptual Framework for Policy-Driven Security

Figure 8 shows NAC's conceptual framework for policy-driven security services. In simple terms, it stores electronic policy representations[9] in a policy repository so that they can be referenced at runtime to make and enforce policy decisions. The following are the principal components of the policy model:

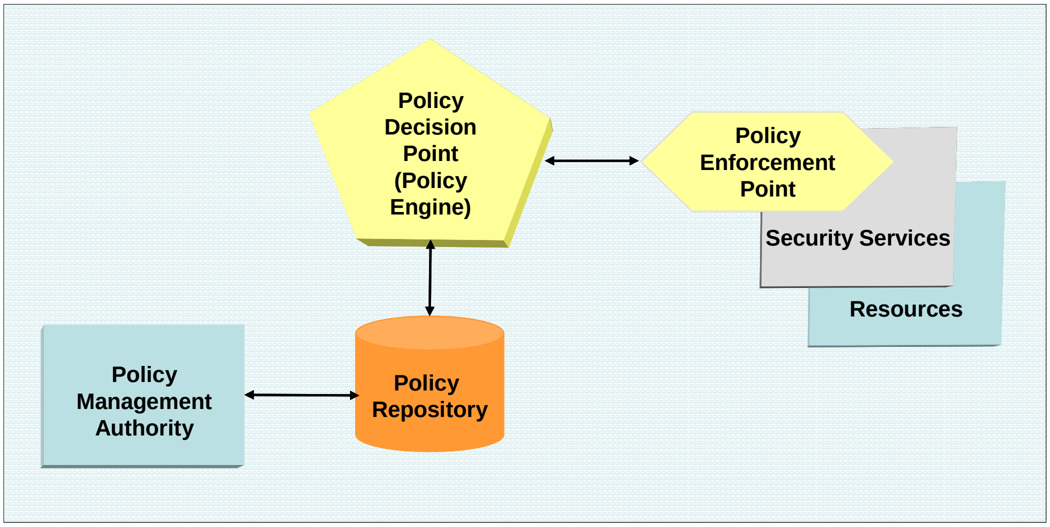

Figure 8. Policy Driven Security Conceptual Framework

- Policy management authority (PMA): the PMA is a person or application entity that composes or creates electronic policy representations through a policy console, policy interpreter, or other tool. These electronic policy representations may be expressed in XML-based policy language, directory entries, configuration file entries, or some other form. Often they are

[9] It should not be assumed that all security policy will be represented electronically. Some policy is of a management nature and will be implemented primarily as management standards. Other policy is of a technical nature and will be implemented primarily as technical standards. It is the latter that will be represented electronically, technology permitting. Policies that coordinate and define the interaction of other technical policies (policies about policy prioritization or conflict resolution) may be difficult to represent electronically, but even these can be addressed by detecting conflicts and ensuring that they are surfaced to the appropriate authority.

configured directly into a proprietary product's policy interface or policy repository. There may be multiple PMAs for different policy domains[10].

- Policy decision point (PDP): the PDP accesses electronic policy information and makes runtime policy decisions at the request of a policy enforcement point (PEP). The PDP is sometimes collocated with the PEP due to product packaging, often justified by performance considerations. In other situations, it is desirable to decouple the PDP from the PEP. There may be multiple PDPs and PEPs, but overall there should be fewer PDPs than PEPs, so as to reduce policy administration and/or allow PDPs to offload complex logic from PEPs.
- Policy repository: the repository is where electronic policy representations are stored. Repositories may be general purpose directory services, or they may be service-specific policy repositories associated, for example, with a specific access management product.
- PEP: the PEP enforces policy at runtime, based on the policy decisions made by the PDP. PEP functionality may be tightly integrated with the security service, as in the case of typical file system, database, or firewall product implementations. Alternatively, it may be a separate agent or plug-in that extends a service implementation to provide policy enforcement, as in the case of Web access management agents or Web server plug-ins that control access to Web pages.
- Security services: these are the core functions of the ESA that cooperate to provide a complete enterprise security services system.
- Resources: the IT assets that security services protect.

[10] For a definition of policy domains and their applications, see the Burton Group's VEN Security Model, as described in: **Securing the Virtual Enterprise Network: Layered Defenses, Coordinated Policies** V2, dated May 23, 2003.

Conceptual Architecture for Policy-Driven Security

With the conceptual framework as a starting point, this section describes NAC's overall conceptual architecture for policy-driven security services. It starts by further decomposing the policy management and security services components of the framework to the specific conceptual services shown in Figure 9. It then describes in further detail the policy decision and enforcement point concepts.

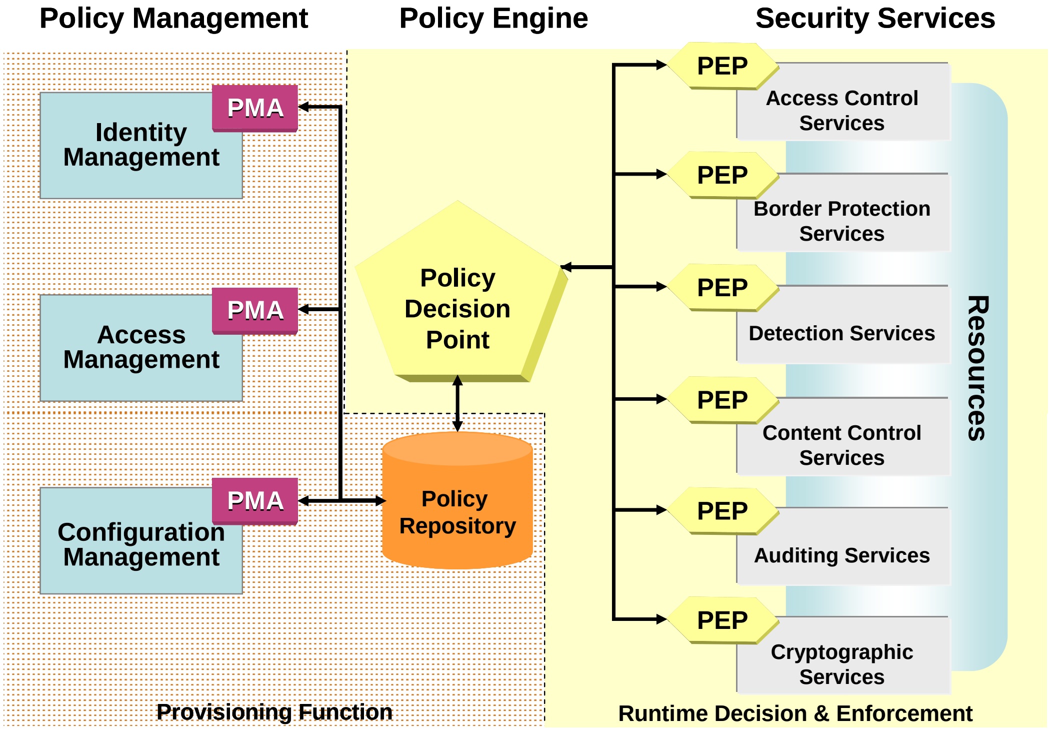

Figure 9. Policy Driven Security Conceptual Architecture

As shown on the left of the figure, policy management has been split into identity management, access management, and configuration management services, which represent three roles of the PMA shown in the conceptual framework. Management services are responsible for maintaining their electronic representation of runtime policy information in the policy repository. A provisioning function can be used to automate the process of updating the policy repository in a timely fashion, so that runtime policy decisions are accurate within an acceptable timeframe. The Identity Management Architecture section addresses provisioning in more detail relative to the creation and maintenance of user accounts for digital identities and their attributes.

On the right are the specific runtime security services and their associated resources and PEPs. The PEP and PDP interact to make runtime policy decisions and then to enforce those decisions via the PEP and the associated service. Again, the level of integration between PEPs and services and between PEPs and PDPs may vary widely.

Following is a brief overview of policy management and runtime security services:

- Identity management services are responsible for assigning and maintaining digital identities and associated attributes across the electronic computing environment and for deleting identities when they no longer represent valid users of the environment.
- Access management services are responsible for assigning and maintaining resource access privileges across the electronic computing environment and for terminating those privileges when they are no longer required. Access management services may encompass a variety of components such as access policy definition, account creation, and access control list (ACL) maintenance. The key differentiator in ESA between access management and identity management is that access management is target-centric or resource-centric, while identity management is initiator-centric or user-centric.
- Configuration management services are responsible for consistently setting and maintaining the security configuration across the electronic computing environment. Configuration management is where NAC's ESA extends the policy-driven conceptual framework beyond the access control framework to include the distributed components of all the security services. Configuration of the various security services—border protection, threat detection, content control, auditing, cryptography, and even configuration management itself—is constrained by the policy model. A classic example is the centralized management and deployment of anti-virus definition files—policies are defined, and updates are automatically pushed to all appropriate corporate endpoints in accordance with that policy.
- Access control services are responsible for controlling access to the enterprise computing environment based on the user's identity (authentication services) and controlling access to specific resources within the environment based on the user's entitlements or privileges (authorization services). This is the classic PDP-PEP implementation where information provided by identity management and access management is used to determine access authorizations.
- Border protection services are responsible for controlling information traffic across external or internal boundaries between security zones, based on the location of the traffic source and destination or on the content of the traffic. In NAC's ESA policy model, configuration of the many devices (including end-user clients) providing border protection services is controlled through centralized policy with configuration definition pushed to the endpoints. Border protection vendors currently provide some tools for centralized management of their proprietary platforms; however, open standards-based, comprehensive management across vendor platforms is generally lacking.
- Detection services are responsible for identifying and protecting against real and potential threats to the computing environment. Policy-based management of detection services generally involves vendor-proprietary solutions for centralized detection engines, with various means for collecting logs from many sources. The consolidated logs are then analyzed by the

detection engine based on pattern and heuristic analysis to identify intrusion attempts.

- Content control services ensure that the enterprise information base is not corrupted and that the external information base being accessed is legitimate and appropriate for business use. Today's anti-virus and anti-spam services are already within the purview of policy-based management controls. One can readily visualize more sophisticated policy-based controls over virus scanning, spam filtering, and content inspection services as well as the emerging enterprise rights management services. Just as organizational roles and job function may be used to determine access privileges, they might also be used to determine the appropriate level of content control.
- Auditing services are responsible for analyzing security logs in support of security investigations, risk assessments, and related activities. The vision for policy-based management is to be able to define auditing requirements in a centralized policy base that is then enforced at the auditing endpoints. This does not seem to be an area of focus by vendors today.
- Cryptographic services are responsible for enabling the confidentiality and integrity of sensitive data and for higher-level digital signature services. The policy-based vision is to be able to define encryption policy for data both in transit and at rest in a central repository, and then apply the policy based on content tags connected directly to the targets. Digital rights management (DRM) technology is beginning to address this requirement, but it is in its infancy and will need several years to mature.

PDP/PEP Detail

Figure 10 provides additional detail on the PDP/PEP portion of the conceptual architecture. Although the model is based on the ISO 10181-3 Access Control Framework, NAC's ESA applies the model to all of the policy management and

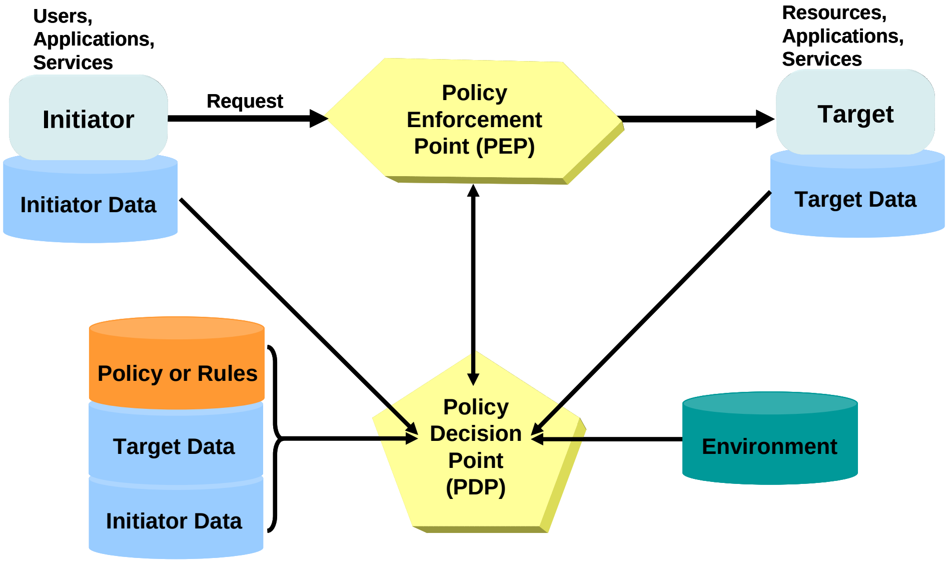
Users, Applications, Services
Initiator
Initiator Data
Request
Policy Enforcement Point (PEP)
Resources, Applications, Services
Target
Target Data
Policy or Rules
Target Data
Initiator Data
Policy Decision Point (PDP)
Environment

Figure 10. PDP/PEP Detail Model

security services that make up the conceptual architecture. The key distinction is that NAC applies the policy decision making and enforcement model to configuration-time services as well as production runtime security services. For additional detail on the importance of this distinction, refer to the vision, technical model, and roadmap for policy automation as described in the Toward Policy Driven Security Architecture section of this document.

As shown above, the PDP/PEP model defines the following key participants:

- Initiator: the user, application, or service that initiates a request of some target resource. The initiator may be a policy management administrator, application, service, or an end user or using application or service.
- Target: the application, service, or other resource to which the request is directed. The target may be a policy repository[11] being updated by a policy management service, or it may be a resource being operated upon by any of the runtime security services.
- PEP: the guard function that enforces policy decisions for target resources.
- PDP: the engine that evaluates requests against the policy (or rules) data and makes policy decisions.

The basic operation is that initiators submit requests to targets. The request specifies an operation to be performed on the target, and it may contain relevant data or more detailed instructions. Requests are intercepted by PEPs, packaged into a decision request, and forwarded to a PDP to determine whether a particular request should be granted or denied. In order to make the policy decision, PDPs may need the following information:

- Initiator data: this is data about the user, application, or service making the request. In the case of a human user, this is known as identity data and could include such things as company affiliation, job function, security clearance, roles, etc. For a service or application, identity data is less clear. In many cases, the service or application is simply a proxy on behalf of some end user, so the identity data will probably be that of the end user. If the application or service is working independently of a specific end user, service identity data might only be a company-defined service ID. In a more sophisticated model, services might be assigned permission attributes. Initiator data is most often stored in an LDAP directory, where the information is created and maintained by identity management services.
- Target data: target data is data about the target resource and is typically related to information sensitivity or content classification.[12] Maintenance and retrieval of target data are among the most difficult functions of policy-driven security architecture. Today, most solutions require embedding some

[11] Policy repository is being used in the broad sense, to encompass any type of policy store, including configuration files residing for example on a general-purpose or special-purpose server or appliance.

[12] Many refer to the assignment of this kind of data as Information or Content Tagging – the act of attributing (tagging) content via metadata to facilitate any or all of the following: information protection (confidentiality, export control classification), information management (identity, version, ownership, valid dates, etc.), or information retrieval (subject/taxonomy, business object type, etc.).

of the PDP logic into the end application or service for dealing with the target data involved in the decision-making process. In the future, a more generalized solution will provide mechanisms for the PDP to query the PEP for target attributes required in the decision-making process.

- Environment data: data about the environment includes details such as time of day, access path, user session context, or transaction context. Access path might indicate the security of the access channel or the current user location (for example, directly connected to the company network or at some Internet café). Environment data such as time of day is easily accessible to the PDP for making decisions. Session context might include strength of user authentication. Transaction context might include the dollar amount of bank withdrawal. Although user location may be very relevant to the decision-making process, under many circumstances that information cannot be reliably obtained.
- Policy or rules: in order to make a policy decision about a request of some resource, we need a statement of the policy (or rule) that can be interpreted by the PDP. In the conceptual architecture, this is shown as the policy repository, which may be a general purpose directory service, a combination of directory data sources accessed as a virtual directory or metadirectory, or a product-specific policy repository.

Once the decision has been made by the PDP, the result would be packaged into a decision reply and returned back to the PEP for enforcement.

Identity Management Architecture

The following sections further analyze two of the identified security services—identity management (IdM) and border protection—to describe service-specific conceptual and logical architectures. These should be viewed as examples that illustrate the logical decomposition of high-level services to the level of detail required to implement the architecture. In terms of the house analogy, they identify the bill of materials required to determine what we need to build or buy. IdM is then discussed at a further level of detail to provide an example of physical architecture, in which the discrete logical services have been mapped to specific products.

For services other than IdM and border protection, only high-level service definitions are provided (see the Other Security Services Template section of the document).

This IdM example first shows the high-level conceptual services, then their decomposition into discrete logical services, and finally their mapping to specific products. Consistent with the overall purpose of the document, these are provided as starting points for developing organization-specific IdM architectures.

Identity Management Conceptual Architecture

Figure 11 depicts the conceptual architecture for identity management (IdM). This is based on the Burton Group's identity management architecture[13], but it is greatly simplified because it focuses solely on the identity administration and provisioning concepts of IdM and does not address access management architecture.

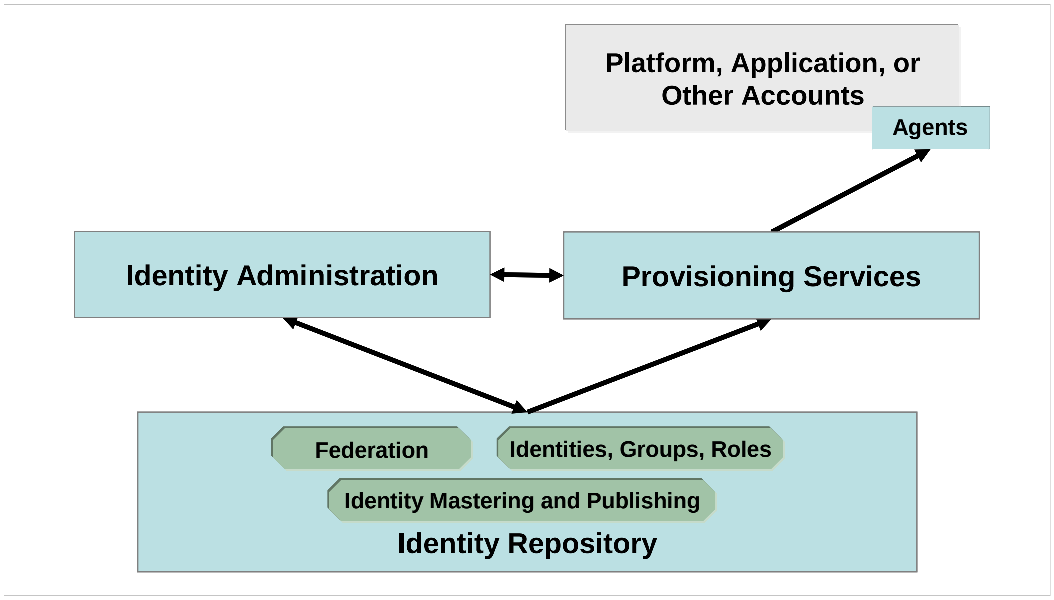
Platform, Application, or Other Accounts
Agents
Identity Administration
Provisioning Services
Federation
Identities, Groups, Roles
Identity Mastering and Publishing
Identity Repository

Figure 11. Identity Management (IdM) Conceptual Architecture

[13] For more complete background information on the topic, see the Burton Group's **Enterprise Identity Management: It's about the Business** v1, dated July 2, 2003.

Following is a brief overview of the key conceptual services of IdM:

- Identity administration services create and maintain unique identities and attributes for various types of users (human users, applications, other digital entities), including external users.
  - It includes delegated administration, self-service administration, and automated administration feeds.
  - It includes identity-mapping services for federated users.
  - Attributes may include roles and groups.
- Provisioning services automate the creation and maintenance of accounts (typically in proprietary systems) through agents associated with supported applications and platforms.
- The identity repository houses identities and their attributes, including federated identities.

Identity Management Logical Architecture

Figure 12 shows the discrete services that make up IdM. Referring back to the house analogy, the goal of logical architecture is to identify the services bill of materials. Thus it should identify the discrete logical services and their relationships at the level required to determine what you need to build or buy in order to construct a set of IdM services for your environment. Note that this is the point in the services

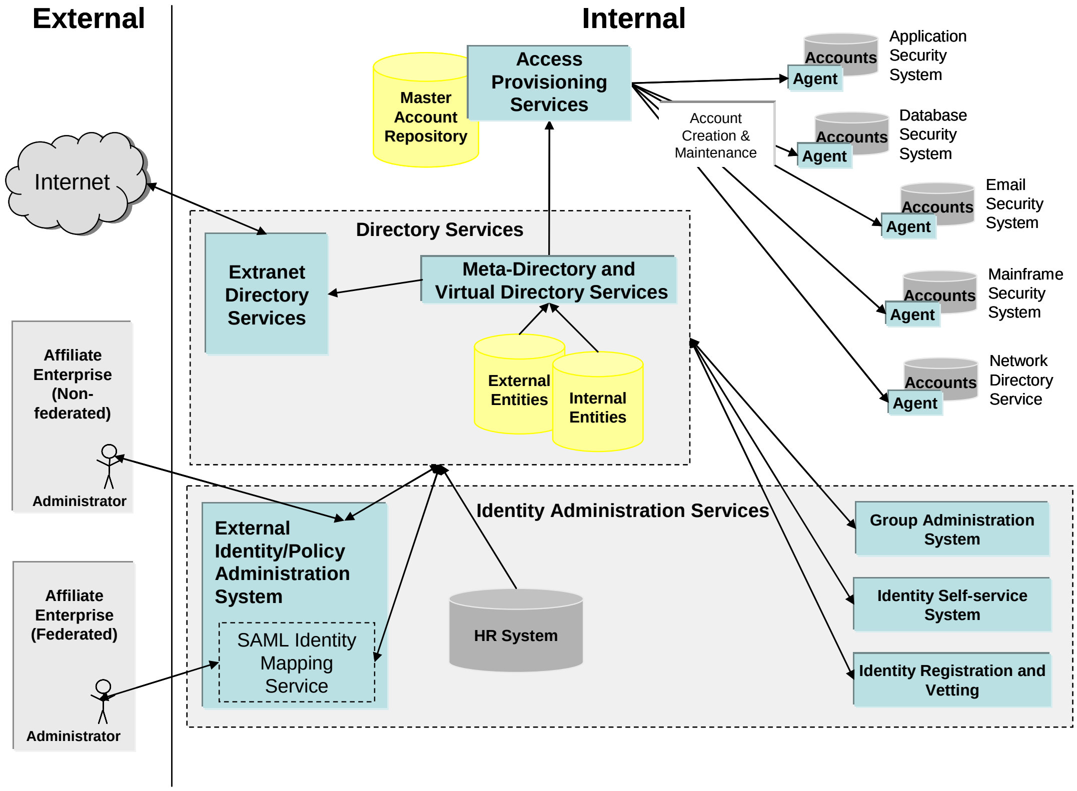
External
Internal
Master Account Repository
Access Provisioning Services
Accounts
Agent
Application Security System
Account Creation & Maintenance
Accounts
Agent
Database Security System
Accounts
Agent
Email Security System
Accounts
Agent
Mainframe Security System
Accounts
Agent
Network Directory Service
Internet
Directory Services
Extranet Directory Services
Meta-Directory and Virtual Directory Services
External Entities
Internal Entities
Affiliate Enterprise (Non-federated)
Administrator
Affiliate Enterprise (Federated)
Administrator
Identity Administration Services
External Identity/Policy Administration System
SAML Identity Mapping Service
HR System
Group Administration System
Identity Self-service System
Identity Registration and Vetting

Figure 12. IdM Logical Architecture

decomposition process where architecture becomes much more organization specific and less generic, so this should be understood as just an example of IdM logical architecture.

The following briefly describes the elements of this IdM logical architecture diagram:

- In the bottom center of the figure is the HR system that provides administrative feeds to create or update internal user identities in the internal entities directory.
- On the bottom left is the external identity administration system that creates or updates external user identities in the external entities directory on behalf of affiliated enterprises (federated or non-federated). In the federated case, security assertions markup language (SAML) identity mapping services are required as well. In both cases, user administration is delegated to an administrator at the affiliate enterprise site.
- In the center are the external and internal user directories, housed in the directory services function. Most enterprises have more than one source of authoritative identity information, including relational databases, mainframe directories, and other LDAP directories. Virtual directory services allow all those sources to be accessed as a single virtual LDAP name space. Alternative solutions such as meta-directories can provide equivalent results. As the Burton Group says, "At the end of the day, product strategies don't matter as much as results. The degree to which an enterprise works to 'clean' its identity house, to 'scrub' the data, to identify authoritative sources, and to make that authoritative data available to key IdM components, will have a huge impact on how successful subsequent IdM efforts will be." Also shown in the directory services component is an extranet directory to provide Internet-accessible directory information to external users. Extranet directories can be used to provide common directory lookup capability or can be a necessary component of scalable inter-company secure e-mail communications and digital signing services.
- Identity registration and vetting functions provide the means for establishing digital identities for persons that might not go through the HR system, such as contractors or consultants. Additional functions may be included to support special identity attributes, such as security clearances or citizenship, which may be provided by organizations other than HR. In some member organizations, a branch of the security organization verifies and maintains these special attributes.
- Identity self-service systems provide for user maintenance of certain identity attributes, as determined by organizational policy.
- Group administration systems allow users to create and maintain groups that provide access control for resources under their control.
- In the top center and right are the provisioning services and agents (not all end systems require agents) that provide account creation and maintenance for the various resource systems. These elements begin to form the identity infrastructure that will be used by access control services.

Identity Management Security Services Template

For completeness, this section provides additional detail on specific IdM services that may be required. These services are responsible for assigning and maintaining digital identities and associated attributes across the environment. This includes deleting or appropriately flagging identities (for historical accountability purposes) when they no longer represent valid users.

User and Identity Administration Services

Identity Administration Services

Identity administration services assign and maintain user and application ("principal") identities and identity attributes, including "federated" identities. The tools typically support centralized and delegated administration of these identities.

Access Provisioning Services

Access provisioning includes those tools and services that maintain access policies and rights.

- Access rules and policies.
- Account and privilege management.

Directory Services

General Purpose Directory Services

- Designed to meet the general needs of many (even unknown) applications.
- Characterized by adherence to international standards and established conventions.
- Provide vendor-agnostic services and are loosely coupled with other infrastructure services.
- Attempt to minimize the need for special-purpose identity stores.
- Include an enterprise directory, which is a general-purpose directory representing the whole population of interest (people, applications, etc.) for the extended enterprise.

Special Purpose Directory Services

- Designed to meet the specialized needs of particular applications or environments such as a network operating system.
- Characterized by vendor-proprietary schemas and features.
- Represent population of interest (users, devices, etc.) to a specific application or environment.

Extranet Directory Services

- Use LDAP proxy and/or LDAP border services to facilitate secure communication and collaboration with business partners.
- Provide a controlled subset of identity information to the public Internet.

- Provide mechanism(s) to obtain directory information from business partners.

Meta-Directory and Virtual Directory Services

- Provide federation capabilities for disparate directory services.
- Provide an abstraction layer between directories and the applications that use them.

Identity Management Physical Architecture

The physical implementation of servers, software, network connections, etc. of the IdM environment described by the logical architecture above is complex. To fully describe such an environment requires multiple documents, including:

- Various diagrams such as software component layering on servers and network topology diagrams.
- Various lists such as all the network addresses of Microsoft Domain Controllers in an environment.
- Documentation of the configuration settings for software components.

This document does not attempt to provide a full set of such documentation because document formats are very specific to individual companies, to the technologies being implemented, and even to the individuals involved. Rather, we provide what should be considered a template for the highest-level view of physical architecture and one from which the need for more detailed documentation can be determined.

Figure 13 maps the discrete services of the IdM logical architecture to specific products, showing their placement on hardware devices and connectivity relationships. It shows how the IdM logical services bill of materials maps to a set of specific products, taking into account that this is remodeling of an existing

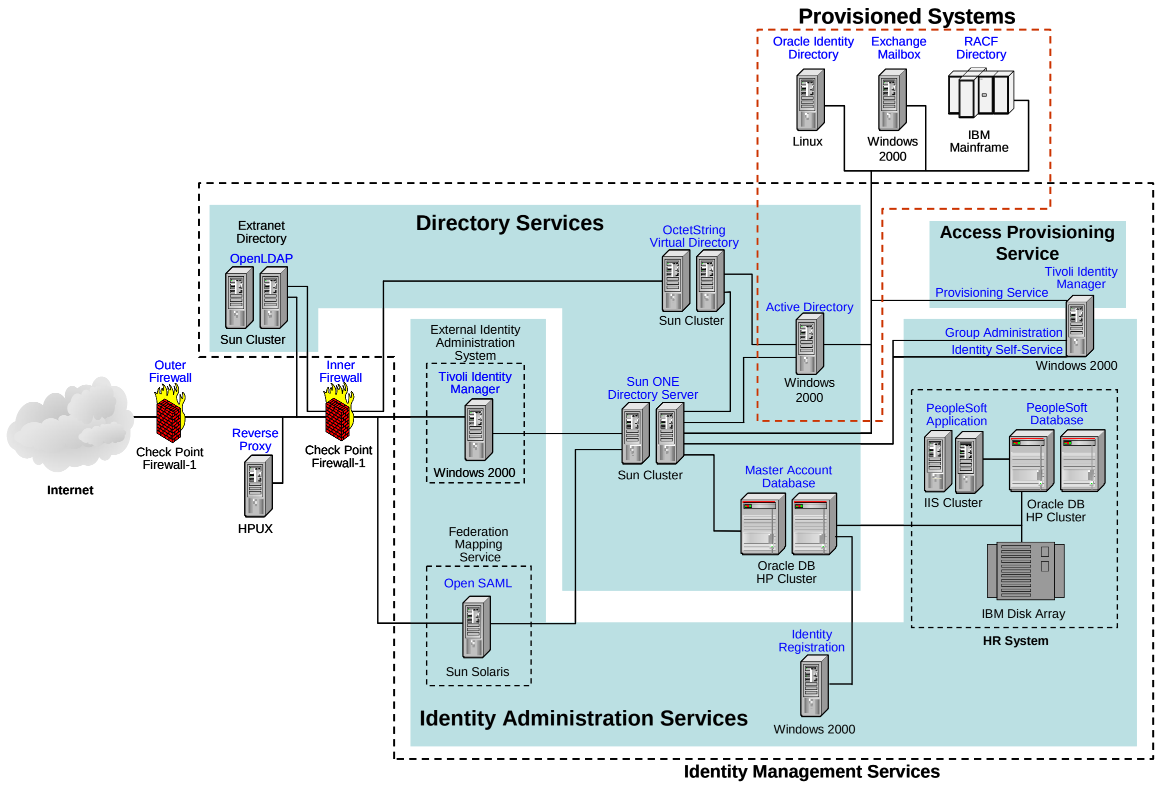

Figure 13. IdM Physical Architecture

infrastructure, not a completely new construction. For example, the existing HR system is a crucial component of the identity administration services.

Some key points need to be understood about the diagram and its use:

- Its purpose is to illustrate a high-level physical architecture diagram, in this case corresponding to the earlier IdM logical architecture, and is not a recommendation on how to structure an IdM environment.
- It is not based on an actual corporate environment or set of requirements.
- No endorsement of specific products is implied. The products listed are intended to be representative of market spaces, and in fact the collection of products shown might be suboptimal for interoperability.
- The diagram shows device connectivity but not information flow. It would be impractical to show all the flows on a single version of the diagram.
- Typically, a separate version of the diagram with a subset of the relevant information flows would be developed for a particular service aspect, such as access provisioning, and used for analysis, communication, and education.

With both the IdM logical and physical architectures now drawn, some key aspects of the interplay between these levels of the architecture hierarchy become apparent:

- The physical architecture, although essential for implementing technology, is much harder to comprehend than the logical architecture and relies heavily on the logical view for context. This fact highlights the need to follow the flow from conceptual, to logical, to physical architecture, not only during design and analysis but also in communication and education of the various audiences touched by the enterprise security program.
- Services that seem distinct at the logical architecture level might be more closely aligned in the physical architecture. For example, in the IdM logical architecture diagram, access provisioning, group administration, identity self-service, and external identity administration services are all distinct. In the physical architecture these services are provided by a single product, Tivoli Identity Manager, and the first three of these services are provided by a single instance of that product.
- Services at the logical level might be decomposed into several sub-services and technologies at the physical level. For example, at the logical level it is adequate to consolidate various directory services into the virtual directory services component. However, to understand the physical architecture it is necessary to expose the five separate directory service implementations in the environment.
- Organizational responsibilities may be far more disjointed than they appear, based on the logical architecture diagram. The administration responsibilities for directory products and identity administration services, for example, may span organizational boundaries.

Border Protection Architecture

This section explains the border protection conceptual and logical architectures by first describing the high-level conceptual services and then providing an example of their decomposition to discrete logical services. Consistent with the purpose of the overall document, these architectures are provided as example templates for use as starting points in developing organization-specific architectures.

Border Protection Conceptual Architecture

One of the key concepts of border protection is that the services are distributed throughout the enterprise; they are not intended to focus only on the boundary between the intranet and the Internet (Figure 14). Another key concept is that packet flow is controlled independently of the initiator’s identity. Thus it might also be referred to as traffic-based access control, as opposed to identity-based access control. This conceptual architecture identifies three primary types of perimeter control. It also addresses controls between the general internal network and isolated “enclaves” and controls on the client platform, as follows:

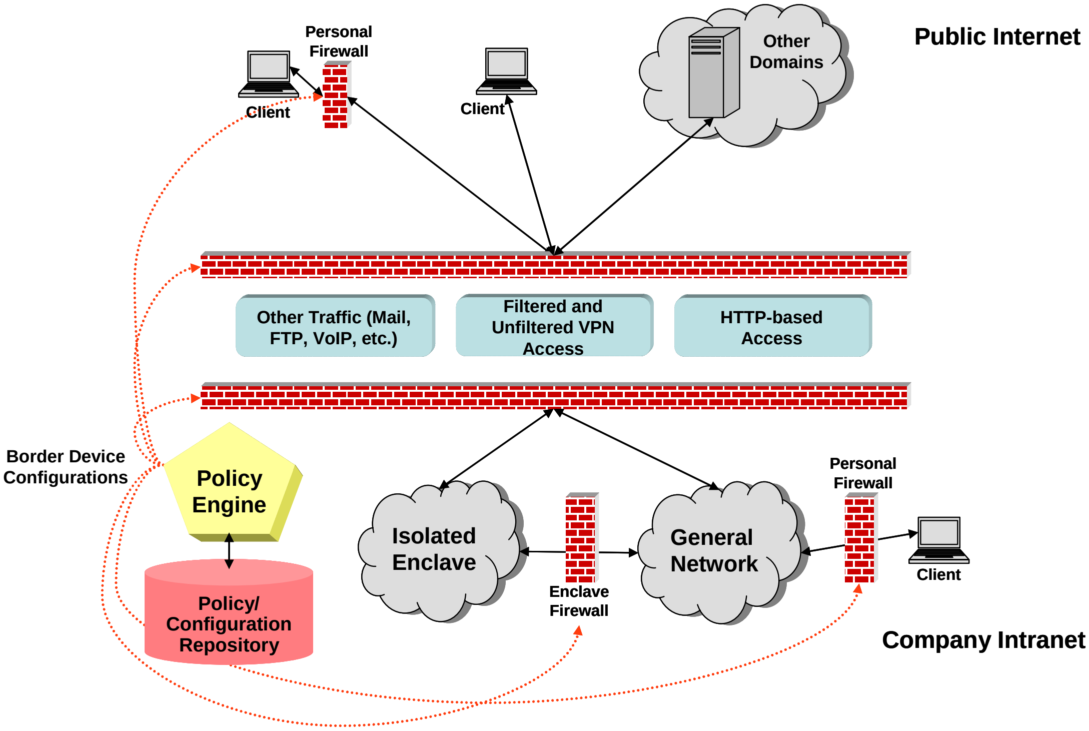

Figure 14. Border Protection Conceptual Architecture

- Filtered and unfiltered virtual private network (VPN) access controls the ability of customers, employees, partners, and suppliers to connect to the company intranet. VPN devices can employ filter lists that restrict incoming access to a specified subset of the applications, services, and other resources inside the company. Where appropriate, unfiltered and unrestricted access can be allowed as well.

- HTTP-based access is the typical means for supporting “e-business.” It allows only HTTP and HTTPS protocols, with access to selected internal company Web resources through reverse proxy technology deployed on the company perimeter.
- The other traffic component provides for the other various types of traffic that must be accommodated, such as e-mail, file transfer protocol (FTP) and voice-over IP (VoIP), which is a rapidly emerging IP telephony technology.
- Enclave firewall: firewalls inside the company are used to isolate special portions of the internal network from the rest of the company intranet. These isolated “enclaves” can be used to either restrict communications from within the enclave out to the general intranet, or to prevent general intranet communications from entering the enclave.
- Client: a client may be a desktop machine (e.g., at home or the office) or a mobile laptop that sometimes connects via the company intranet and other times connects via the public Internet (at home or on the road). In other cases, the client may simply represent a client service.
- Server: the server box symbol shown in “Other Domains” represents other external services that may require access to the company intranet, or external services that require access from the intranet.
- Personal firewalls are deployed to client machines to prevent unauthorized communication to the client, as well as protecting clients from worms and other invasions that evade detection by anti-virus software—both when connected to the company intranet and when connected directly to the public Internet. For example, a personal firewall would prevent unauthorized access to the machine while it is connected via an Internet café. This model shows two client machines with personal firewalls, one inside the company perimeter and one in the public Internet. The external client uses a VPN connection to get back into the company intranet to protect the traffic between the company perimeter firewall and the client personal firewall.

Border Protection Logical Architecture

Figure 15 decomposes the border protection conceptual architecture to identify the discrete logical services and their relationships at the level required to determine what you need to build or buy in order to construct border protection services for your environment. Again, this is the point in the process where architecture becomes much more organization specific and less generic, so this should be understood as an example of border protection logical architecture.

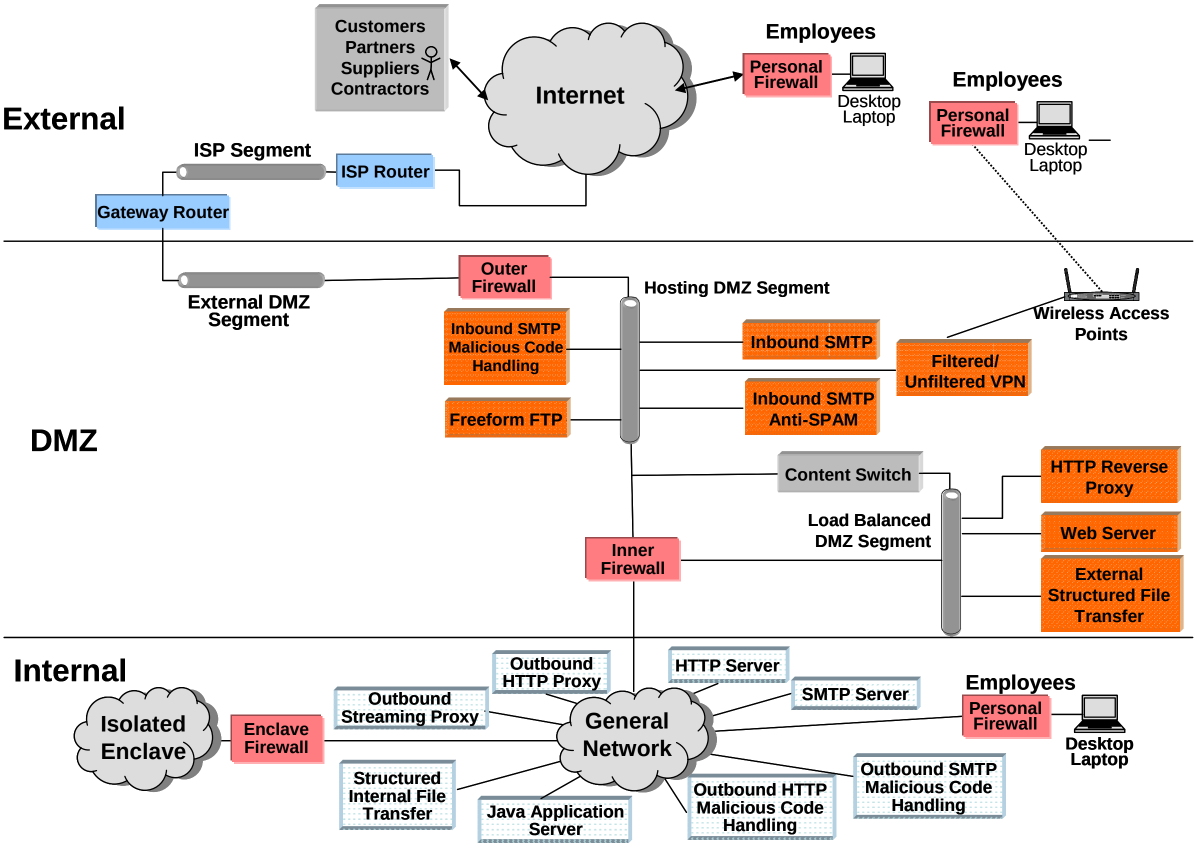

Figure 15. Border Protection Logical Architecture

The following briefly describes the primary elements of the logical architecture that were not covered in the conceptual architecture description:

- The gateway router manages Internet routing and provides coarse-grain packet filtering based on IP/TCP/UDP protocols.
- The external demilitarized zone (DMZ) segment is a limited functionality network segment that provides connectivity between the gateway router and the outer firewall
- The outer firewall is a high-performance, low-latency firewall capable of providing:
  - An additional layer of IP/TCP/UDP packet filtering.
  - In-depth packet inspection and protocol validity checking.

    - Some level of denial of service (DoS) detection and prevention.
    - Secured IP routing to mitigate IP address space leakage.
- The hosting DMZ segment is a network segment between the outer and inner firewalls that may contain hosts/servers based on the need to locate them behind a load-balancing content switch
- Wireless access points (APs or WAPs) are transceivers in a wireless LAN that act as transfer points between wired and wireless signals, and vice versa. An AP is sufficiently trusted to put it inside the outermost firewall, so it's marginally more trustworthy than the Internet at large. However, its ability to implement security (namely authentication, authorization, and encryption) is deemed insufficient and a VPN[14] is used to implement those functions on top of the wireless infrastructure (the same way VPN is used to provide secure communication paths over the Internet).
- The content switch is a traditional IP load-balancing device that also has the capability to balance sessions (TCP) across servers. It is a key mechanism for hiding the internal IP address space.
- The load-balanced DMZ segment is similar to the hosting DMZ segment except that it is behind the content switch and provides a network segment between the content switch and the inner firewall.
- The inner firewall is a high-performance, low-latency firewall capable of providing:
    - An additional layer of IP/TCP/UDP packet filtering.
    - In-depth packet inspection and protocol validity checking.
    - Some level of DoS detection and prevention.
    - Limited, secured IP routing or, more often, static IP routes that mitigate IP address space leakage or unauthorized IP traffic.

Border Protection Security Services Template

Border protection services control information traffic across external or internal boundaries between security zones, based either on the location of the traffic source and destination or on the content of the traffic. In NAC's ESA policy model, configuration of the many devices (including possibly end-user clients) providing border protection services is controlled through centralized policy with configuration definition pushed to the endpoints. Border protection vendors currently provide some tools for centralized management of their proprietary platforms; however, open standards-based, comprehensive management across multiple vendor platforms is generally lacking.

[14] An alternative is to use the Extensible Authentication Protocol (EAP), possibly in conjunction with proprietary vender features, to sufficiently secure the wireless infrastructure and associated endpoints. EAP is a general protocol for authentication that also supports multiple authentication methods, such as token cards, Kerberos, one-time passwords, certificates, public key authentication and smart cards. IEEE 802.1x specifies how EAP should be encapsulated in LAN frames.

Packet Filtering Service

- This service provides dynamic, stateful (i.e., stable), IP-only packet filtering. It examines each IP packet and determines whether to allow the packet to pass. It does this by examining source and destination addresses and ports and by understanding the "state" of a transaction. This allows a reply to be treated differently than a query.
- Packet filtering also supports the concept of one service per server and segregation of servers from each other. Packet filters deny all traffic to the server that is not expressly allowed; for example, an HTTP proxy receives only HTTP requests, and a VPN server receives only VPN requests.
- Packet filtering provides security controls between different zones of trust.

VPN Service

A VPN encrypts data as it traverses un-trusted networks. VPNs fall into two categories: LAN-to-LAN and client-to-server. The main difference is that client-to-server VPNs usually require a user to authenticate (e.g., by providing a user name and password), whereas LAN-to-LAN VPNs do not. LAN-to-LAN VPN "tunnels" are usually set up between routers or servers and are transparent to users.

Proxy Services

Forward Proxy Services

These services provide access to the external Web while protecting corporate workstations from external threats and filtering offensive materials from coming into the corporation. The services support HTTP, HTTPS, and FTP protocols and are outbound only, so that requests must be initiated from inside the corporate network.

Reverse Proxy Services

These services enable secure external access to internal corporate resources by requiring user authentication and authorizing user access only to selected locations on interior Web servers. The reverse proxy service transmits authorized requests to the permitted interior server and then returns the response from the interior server to the user.

Application Proxy Services

These services provide inbound and outbound connection between the Internet and the corporate intranet in support of FTP, Telnet, TN3270, SQLNet, X-Windows, and line-printer daemon (LPD) protocols.

Other Security Services Template

The preceding Identity Management Architecture and Border Protection Architecture sections provided example architecture diagrams and descriptions for those services. This section discusses each of the other security services identified in the

Conceptual Architecture for Policy-Driven Security on page 45. These services are described in some depth and broken out into second-level and in some cases third-level services. These definitions are based on input from several member organizations and the Burton Group and are intended to serve as a template that organizations may choose from and tailor to their specific current and future needs.

Access Management Services

Access management is responsible for assigning and maintaining resource access privileges across the electronic computing environment and for terminating access privileges when they are no longer required. Access management services may encompass a variety of components such as access policy definition, account creation, and ACL maintenance. The key differentiator in ESA between access management and identity management is that access management is target-centric or resource-centric, while identity management is initiator-centric or user-centric.

Configuration Management Services

Configuration management is responsible for consistently setting and maintaining the security policy configuration across the electronic computing environment. As discussed earlier, this is where NAC's ESA extends the PDP/PEP model to support configuration-time instantiation of policy for all management and security services. This includes instantiation of policy decision and enforcement data for the identity, access, and configuration management services themselves, and all the various production runtime security services—access control, border protection, threat detection, content control, auditing, and cryptography. A current example is the centralized management and deployment of anti-virus definition files—policies are centrally defined, and updates are automatically pushed to all affected corporate endpoints in accordance with that policy. A more forward-looking example is the centralized management of policy configuration files for border protection and threat detection servers, which eliminates the need to manually configure each end point based on configuration checklists for each variant of the server architecture.

Access Control Services

Access control services are responsible for controlling user access to the enterprise computing environment based on the user's identity (authentication services) and controlling access to specific resources within the environment based on the user's entitlements or privileges (authorization services). This is the classic PDP-PEP implementation where information provided by identity management and access management is used to determine access authorizations.

Authentication Services

In simple terms, authentication services verify *who* the user is. Typically, the user is required to have a unique identity, and that identity is then verified using a particular authentication technique. It may be a human user with a unique name or user ID, or a computer process with a unique ID. Authentication services generally also assign the user to one or more roles or groups with identities that are subsequently used for authorization. Authentication services are the linchpin of user security—everything else depends on knowing the unique identity of the user and the roles or groups to which it belongs.

Direct (First-Person) Authentication Services

Direct authentication services verify the unique identity of the human user or process based on a unique user identity and password or a stronger authentication technique (smartcards, secure ID fobs, biometrics, etc.).

Indirect (Third-Party) Authentication Services

Third-party authentication services are trusted services that pass previously authenticated identities. They may include single sign-on (SSO) products, SAML-based services, perimeter proxies, etc.

Authorization Services

Authorization services determine *what* a properly authenticated user is allowed to do. The authorization process ensures that users are allowed only the access they require to do their jobs. This is referred to as the principle of least privilege and is as important for electronic users (processes or applications) as it is for human users. As an example, adherence to the principle of least privilege in program design reduces the damage that can occur if a user attempts to exploit that program for mischievous or malicious purposes.

Online (Connected) Authorization Services

Online authorization services evaluate identity, environment, and asset attributes (tags) against policies (sets of rules) to arrive at a permission recommendation. This process includes the use of distributed, on-demand authorization services.

Offline Authorization Services

Offline authorization services correctly determine permission in offline situations and in potentially hostile environments. This is a future direction that may not be realized in the next 3 to 5 years.

Detection Services

Detection services assist in the enforcement of security policy through the ongoing creation, capture, and monitoring of security-relevant events. The goal is to detect and respond to threats and vulnerabilities in a way that prevents damage or loss. One of the design requirements[15] is to gracefully degrade access to services in the event of an attack or disaster, to recover from the resulting failures, and to efficiently restore access as impeding circumstances wane.

[15] The goal is a resilient design that adapts to attacks or disasters in reasonable ways. Resilience is the property often associated with disaster recovery processes, defenses against denial of service attacks, fallback regimes used for restoration of critical services, and similar approaches to assuring availability.

Policy-based management of detection services generally involves vendor-proprietary solutions for centralized detection engines, with various means for collecting logs from many sources. The consolidated logs are then analyzed by the detection engine based on pattern and heuristic analysis to identify intrusion attempts.

Intrusion Detection Services

Intrusion detection services identify attempts to break in to a protected network or system and provide real-time or near real-time alarms.

Anomaly Detection Services

Anomaly detection services identify irregularities or glitches in the infrastructure that could be exploited, and they provide guidance for dealing with the risk.

Vulnerability Assessment Services

Vulnerability assessment services are used to analyze systems to identify potential security weaknesses and exposure to known threats.

Logging Services

Logging services provide the capability to collect and consolidate security logs.

Content Control Services

Content control services ensure that the internal enterprise information base is not corrupted and that the external information base being accessed is legitimate and appropriate for business use. In addition, enterprise digital rights management technology is now being introduced to provide content-based control over what can and cannot be done with information.

Today’s anti-virus and anti-spam services are already within the purview of policy-based management controls. One can readily visualize more sophisticated policy-based controls over virus scanning, spam filtering, and content inspection services as well as the emerging enterprise rights management services. Just as organizational roles and job function may be used to determine access privileges, these attributes might also be used to determine the appropriate level of content control.

Anti-Virus Services

Anti-virus services identify, block, and remove viruses embedded in e-mail and files. Major virus targets include boot records, program files, and data files with macro capabilities (e.g., Microsoft Word document and template files). Viruses spread rapidly as infected program and document files are shared via e-mail, and they are also transmitted through direct downloading from Internet sites and through sharing of removable media that are infected.

Anti-Spam services

Anti-spam services attempt to identify spam e-mail messages and filter them out. An alternative strategy that may be appropriate in some environments is to filter out the good e-mail and assume that everything else is spam.

Enterprise Rights Management Services

Enterprise rights management (ERM) services utilize digital rights management technology to govern access to enterprise information throughout its life cycle. Traditionally, digital rights management has been used commercially to protect electronic media such as music and movies. Recently, it is being used to protect sensitive information both inside and outside the enterprise. ERM software provides fine-grain control over what can and cannot be done with information. For example, e-mail messages can be marked with usage permissions and identity-specific access controls so that they can be neither modified nor forwarded to parties outside the organization. For information that once was very difficult to protect, such as word processor documents on removable media, rights management provides policy-based cryptographic protection even when the data is physically stolen.

Content Inspection Services

Content inspection services utilize content inspection technologies to detect and then deal with viruses, spam, and pornography or other information content control issues.

Auditing Services

Auditing services are responsible for aggregating, normalizing, and analyzing events from consolidated security logs in support of day-to-day event management and other security-related activities. Audits may be conducted to ensure the integrity of information resources, to investigate incidents, to ensure conformance to security policies, or to monitor user or system activity as appropriate.

The vision for policy-based management is to be able to define auditing requirements in a centralized policy base that is then enforced at the auditing endpoints. This does not seem to be an area of focus by vendors today.

Cryptographic Services

Cryptographic services enable the confidentiality and integrity of sensitive data, and provide higher-level digital signature services. Enabling services are designed to handle the details of cryptographic key management and support on behalf of using services. Higher-level digital signature services can be used to authenticate the identity of the sender of a message or the signer of a document and to ensure that the original content of the message or document is unchanged.

The policy-based vision is to be able to define encryption policy for data both in transit and at rest in a central repository and then apply the policy based on content tags connected directly to the targets. DRM technology is beginning to address this requirement, but it is in its infancy and will need several years to mature.

Cryptography Services

- Enabling services implement standard cryptographic algorithms on memory objects, documents, files, repositories, data streams, etc.

- Some of the cryptographic topics supported include encryption, hashing, key generation, digital watermarking, and steganography.
- These services may be delivered by various servers, Web services, and desktop tools, but primarily developer libraries.

Public Key Infrastructure Services

- Public key infrastructure services manage and process X.509 V3 certificates, including the certificate authority, the certificate revocation list, certificate validation services, and trust relationships.
- This definition specifically excludes identity management and key generation.

Private Key Storage Services

These trusted services store private keys, including key escrow for private and secret encryption keys, personal key wallets, smart cards, and hardware key vaults for private signing keys.

Digital Signature Services

Digital signature services can be used to authenticate the identity of the sender of a message or the signer of a document and to ensure that the original content of the message or document is unchanged. Digital signatures are easily transportable, cannot be imitated by someone else, and can be automatically time-stamped. They provide the capability to ensure that the original signed message arrived, which means that the sender cannot easily repudiate it later. Assuming the required legal status exists, they can enable support of non-repudiation requirements.

Signing Services

Signing services provide the capability to sign documents, files, and memory objects electronically, on behalf of an identity. Implementation may include desktop signing tools, a signing server (server-hosted signing keys for browser-only signing), various workflow applications, and developer libraries.

Notary Services

Notary services provide trusted long-lived digital signatures and timestamps on top of existing, valid signatures.

Code Signing Services

These are services used to sign code and other programming deliverables.

Verification Services

These services are used to verify signatures and establish data integrity.

This ends the description of the Security Technology Architecture components and services started on page 42. The next section is focused on design and development process.

Design and Development

This section identifies the types of guidance that organizations may want to provide to those responsible for design, development, and deployment of applications. This guidance applies to enterprise security infrastructure components as well as applications that use the infrastructure. Further, much of this guidance can be applied to components or applications built in-house as well as commercial off-the-shelf (COTS) applications selected for integration. The goal is to select applications that utilize their ESA services in the most effective way possible, so as to achieve the overall security goal[16] and objectives[17] for the organization.

Design and development guidance may range from overall process guidelines to specific guides, templates, and tools. It may include design patterns, code samples, reusable libraries, and testing tools. All of these are aimed at effective utilization of ESA and effective integration into the ESA environment.

The following discussion is based on NAC member organization experience and is intended to serve as a starting point for an overall process outline, with a few notes about each element and in some cases references to additional information.

Design Principles

Once the process of *identifying* the guiding principles has been completed, as described in the Policy Framework Overview starting on page 14, then those guiding principles are used as security design principles. They may be tailored or augmented to make them more design and development specific, and they then become the starting point for designing and developing ESA applications. A design principles checklist should be provided to all those responsible for design, development, and testing of ESA applications.

Design Requirements

Design requirements can be categorized as explicit or implicit. However, to support Requirements-Based Testing, all requirements need to be made explicit as part of the specifications process. A requirements checklist should be provided to all those responsible for design, development, and testing of particular ESA applications.

Explicit Requirements

Business Requirements

These are new business opportunity requirements.

[16] The IT security goal is to enable an organization to meet all mission/business objectives by implementing systems with due care consideration of IT-related risks to the organization, its partners, and its customers.

[17] The five security objectives are availability, integrity, confidentiality, accountability, and assurance.

Compliance Requirements

These include legal, legislative and regulatory requirements (e.g., HIPAA, Sarbanes-Oxley).

Technology/Deployment Policy Requirements

This category covers infrastructure requirements, as opposed to the business requirements named above. As an example, if there are server-to-server authentication and connectivity requirements, they could affect the application design in some way.

Implicit Requirements

Data Class

These requirements are based on the confidentiality or privacy classification of the data (e.g., if a Social Security number is passed over the wire, the risk of passing it over a particular type of data channel must be assessed).

Threats

These requirements are based on known threats in the particular application environment.

Design Best Practices

Design Patterns

Design patterns are recurring solutions to software design problems that are ubiquitous in real-world application development. Patterns give information system architects a method for defining reusable solutions to design problems in a way that is programming-language independent. Following are some references for security design patterns:

- Architectural Patterns for Enabling Application Security, by Joseph Yoder and Jeffrey Barcalow (Copyright 1998)
- Security Design Patterns Part 1 v1.4, by Sasha Romanosky, November 12, 2001
- *Security Design Patterns*, by Bob Blakley, Craig Heath and members of The Open Group Security Forum, Copyright April 2004, The Open Group. The following link will take you to The Open Group Web site where you can access a free PDF version on the Web: The Open Group Security Design Patterns.

Security Engineering

Security engineering is the field of systems engineering dealing with the security and integrity of real-world systems. It is the engineering discipline focused on how to build dependably secure systems in the face of malice, error or mischance. It requires expertise in a variety of disciplines – including computer security, cryptography, applied psychology, management and the law – as well as knowledge of critical applications. The following book on security engineering was recommended by Fred Cohen of the Burton Group and has rave reviews at Amazon.com:

- Security Engineering: A Guide to Building Dependable Distributed Systems, by Ross J. Anderson (Wiley, 2001)

Ross Anderson is one of the pioneers of security engineering as a formal field of study at Cambridge University.

Reusable Tools, Libraries, and Templates

To ensure proper utilization of the security infrastructure and to simplify the job of the developers and system administrators, it is important to provide meaningful guidance at the code level. These objectives require detailed guidance that is specific to each platform (i.e., hardware/software combination) endorsed by the enterprise.

This is a critical step in an organization’s security architecture development that is easy to overlook. The security architecture is developed from the top down and is typically delivered by those looking at the big picture vision for the enterprise. However, that vision, though necessary, is of little interest to most developers and administrators. The architecture must bridge the gap from a vision, prolific with pictures, to the code and configuration level.

The developer guidance should include:

- Detailed, platform-specific instructions for utilizing the security services. These instructions should include installation and configuration of the required software components for the particular environment, as well as methods to accomplish various security functions such as authentication, authorization, and encryption.
- Sample configurations and working code that is well documented for review by developers.
- Code snippets that can be freely copied or downloaded for incorporation into development projects.
- Libraries that provide necessary security functions that have encapsulated low-level code details into simpler, higher-level functions.
- Templates that describe typical application functionalities with necessary security aspects identified.

Coding Best Practices

Design is complete; design patterns have been identified; security engineering principles have been taken into account; and reusable tools, libraries, and templates have been put in place. Now secure design must be translated into secure code. The starting point is to understand the best practices for developing and delivering secure code and then to utilize a process that supports those practices. The following book had received rave reviews at O'Reilly.com and Amazon.com:

- *Secure Coding: Principles and Practices*, by Mark G. Graff and Kenneth R. Van Wyk (O'Reilly, 2003)

Two larger books that also have good reviews are:

- Building Secure Software: How to Avoid Security Problems the Right Way, by John Viega and Gary McGraw (Addison-Wesley, 2001)
- *Writing Secure Code, Second Edition*, by Michael Howard and David C. LeBlanc (Microsoft Press, 2002)

Code Reviews

Your organization may want to consider putting a code review process in place if it hasn't already. The book, *Peer Reviews in Software: A Practical Guide,* by Karl E. Wiegers (Addison-Wesley, 2001) has excellent reviews at Amazon.com.

Code reviews focus on more than just security issues, but one of their key purposes should be to review for secure coding practices.

Input Validation

One of the first lines of defense in any secure program is to validate input. The article, "Best practices for accepting user data," provides information on that topic.

In addition, the following excerpt from a June 4, 2003, *Computerworld* article covers a subset of the same issues and seems to be consistent in its recommendations: [The title and author of the article should be noted here, and quotation marks need to be added to indicate the beginning and end of the excerpt.]

Application-specific vulnerabilities are born in coding primarily because developers fail to validate user input. Not detecting and eliminating this simple error can allow the following exploits to occur:

- Buffer overrun: A buffer overrun occurs when data larger than the entry field is written into memory. Hackers can use this flaw to overload the server with data and crash the site, shutting down business.
- Cross-site scripting: Cross-site scripting happens when a hacker injects malicious code into a site to make a user session appear as if it is originating from the targeted site. As a result, the attacker is given full access to any information exchanged in the user session, such as account passwords and Social Security numbers.
- SQL injection: An SQL injection embeds SQL script into the user input by placing a malicious character (such as an apostrophe) in the input field, allowing the SQL server to execute a malicious query such as delivering a directory of customer credit card numbers.

Preventing these attacks primarily requires a change of mind-set. Simply rewriting a few lines of code to apply the following maxims will greatly reduce the risk of attack:

- Validate all user input to allow nothing other than the function's expected input and output.

- Set a trusted code base, and validate all data before it can enter the trusted environment.
- Test each data type before entry (i.e., Web, registry, file system, and configuration files).
- Define all data format (such as buffer length and integer type).
- Define valid user requests and reject all others.
- Look for valid, not invalid, data.
- Never mirror Web input.
- Encode or encrypt output.

Code Analysis Tools

Source code analysis products are aimed at helping companies unearth and fix flaws in C/C++ and Java code-based application development. The goal is to give companies a way to discover flaws in code that could lead to threats such as buffer overflows, format string errors, and SQL injection exploits. They also include a runtime analysis component that allows security workers to launch a variety of attacks against new applications before they are deployed. One of the vendors is Fortify Software, highlighted in an April 05, 2004, *Computerworld* article.

Testing Best Practices

According to the International Institute for Software Testing (IIST), “Experience reports show that up to 80% of the maintenance effort is spent to fix problems resulting from requirements errors.… Well-understood, well-defined and managed requirements are the basis for effective testing of the software system.”

Requirements-Based Testing

Testing should tie back to the *requirements*. As mentioned earlier under Design Requirements, the specifications should make all requirements explicit, including both the *positive* (should happen) and *negative*[18] (should not happen) requirements. Some requirements will be application specific, while others will be general requirements derived from the Design Principles. There should be a *test checklist* for both.

Requirements-Based Testing Tools

Requirements-based testing methodology tools are available. However there is very little related to security testing specifically.

[18] For example, based on design principles any component that controls access to resources should be tested to ensure that it does not fail open (i.e., it fails in such a way that no access is granted).

Security Operations

The focus now shifts to security operations, highlighted in the bottom center of the overall framework in Figure 16. This is the third and final set of components and processes that make up the NAC's ESA.

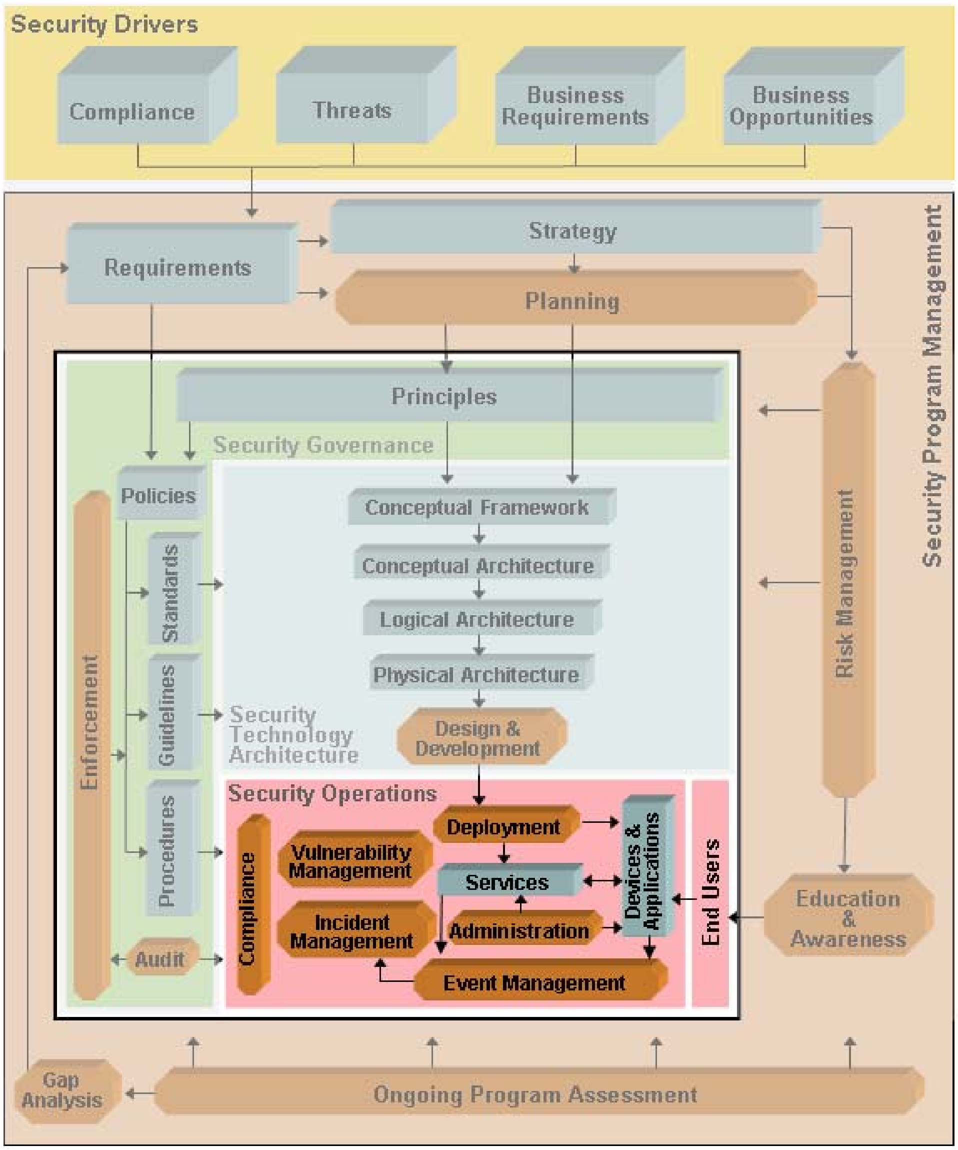

Figure 16. Security Operations Components and Processes

The purpose of this section is to provide an overall, but high-level security operations framework and template that member organizations can tailor to their needs. The security operations function defines the processes required for operational support of a policy-driven security environment. This function involves two key types of processes. One includes the administration, compliance, and vulnerability management processes required to ensure that the technology as deployed conforms to policy and provides adequate protection to control the level of risk to the environment. The other type consists of the administration, event, and incident management processes required to enforce policy within the environment. The security operations function has a strong dependency on asset management. Figure 17 shows this overall set of components and processes.

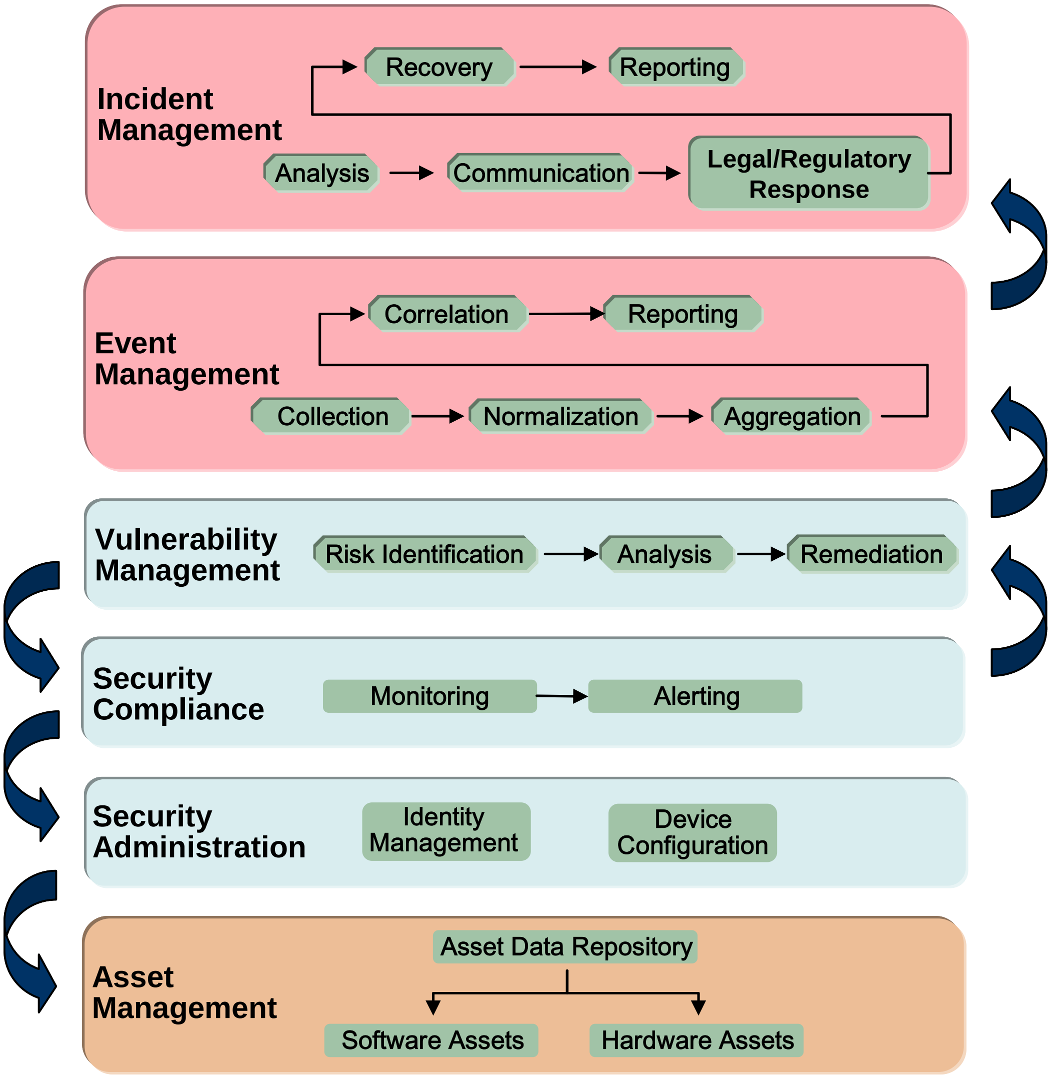

Figure 17. Security Operations Overview

The components and processes that make up security operations are introduced briefly below and then described in more detail in the following sections:

- Asset management is a component and process for maintaining the inventory of hardware and software assets required to support device administration, compliance monitoring, vulnerability scanning, and other elements of security operations. Though not strictly an ESA component, it is a key dependency of security operations.
- Administration is the process for securing the organization's operational digital assets against accidental or unauthorized modification or disclosure.

- Compliance is the process for ensuring that the deployed technology conforms to the organization's policies, procedures, and architecture.
- Vulnerability management is the process for identifying high-risk infrastructure components, assessing their vulnerabilities, and taking the appropriate actions to control the level of risk to the operational environment.
- Event management is the process for day-to-day management of the security-related events generated by a variety of devices across the operational environment, including security, network, storage, and host devices.
- Incident management is the process for responding to security-related events that indicate a violation or imminent threat of violation of security policy (i.e., the organization is under attack or has suffered a loss).

Asset Management

Asset management includes the components and processes for maintaining the inventory of hardware and software assets required to support device administration, compliance monitoring, vulnerability scanning, and other aspects of security operations. Common components include a repository of hardware and software assets (including the configuration and usage information), a capability to discover assets as they are added to the network, and reporting capabilities. This information may be used for activities such as contract renewals, software license compliance audits, and cost reduction activities. While asset management is not specific to ESA and may in fact be valued more for its contribution to enterprise architecture, it is a foundational dependency of security operations.

From a security perspective, an organization must be able to respond quickly to threats, and doing so requires knowledge of the assets that may be under attack.

Asset information needs include:

- Asset location
- Configuration of software and hardware
- Support ownership
- Business context
- Identity information

Monitoring is also required to ensure that the inventory is complete and up to date.

Security Administration

Security administration includes the components and processes for securing the organization's operational digital assets against accidental or unauthorized modification or disclosure. This is accomplished by planning, coordinating, and implementing the technologies and best practices required to create and maintain secure access to resources and protect the integrity of system and device configurations.

Security administration comprises two primary sub-components:

- Identity management is responsible for the creation, modification, and termination (inactivation or deletion) of digital identities, including the workflow process for managing both identity and access management information. It is also responsible for management of authentication tokens and certificates.
- Device configuration is responsible for technical standards instantiation at the device level (see the Policy Automation Vision section of this document for background information). It is also responsible for ensuring that updates to the actual devices are reflected in the asset database.

Security Compliance

Security compliance (Figure 18) provides a process framework for ensuring that the deployed technology conforms to the organization's technical standards, procedures, and architecture. An organization must have processes and tools that enable:

- Monitoring of the deployed technology to ensure that it remains in alignment with policy. Monitoring involves gathering data about deployed technology and comparing it to a defined state.
- Alerting and reacting to identified exceptions, bringing the technology back into alignment. When technology is determined to be out of alignment, processes need to be in place that allow for notification of the appropriate personnel and bringing the technology back into compliance.

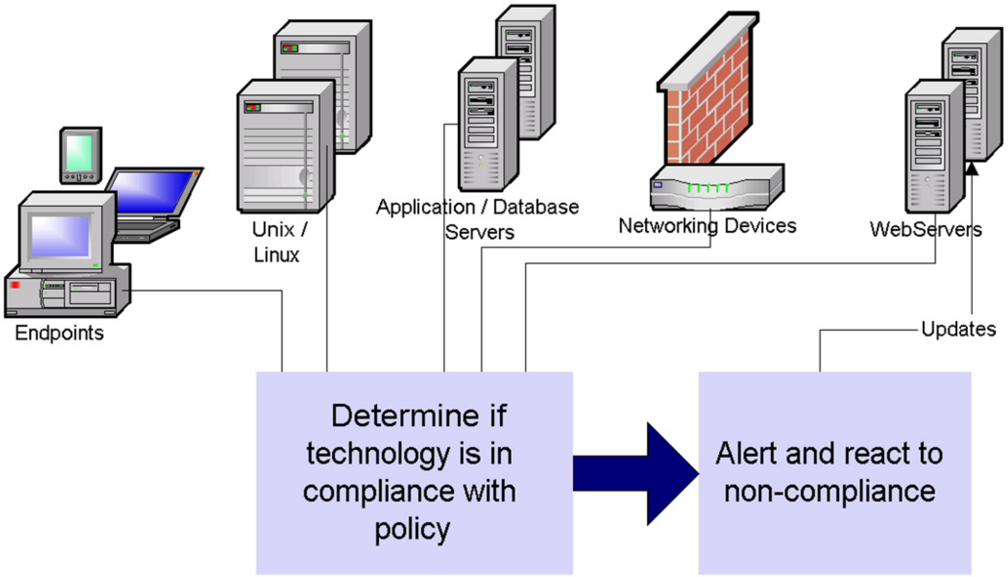
Endpoints
Unix / Linux
Application / Database Servers
Networking Devices
WebServers
Updates
Determine if technology is in compliance with policy
Alert and react to non-compliance

Figure 18. Security Compliance

Vulnerability Management

Vulnerability management provides a process framework for identifying high-risk infrastructure components, assessing their vulnerabilities, and taking the appropriate actions to control the level of risk to the operational environment.

Asset management is a core dependency for the vulnerability management process. It is assumed that the asset repository:

- Contains hardware and software configuration information, owner information, and business context and value information.
- Allows exact understanding of targets for remediation of vulnerability notifications from vendors.
- Supports evaluation of targets identified as a result of vulnerability assessment scanning.

Vulnerability management encompasses both reactive and proactive processes for dealing with vulnerability issues:

- Reactive process for dealing with vulnerability reports from vendors.
- Proactive process for identifying vulnerabilities and taking appropriate actions to control the level of risk.

Reactive Process for Responding to Vulnerability Notifications

- Receive notification of potential vulnerabilities.
- Query the asset repository looking for target systems that are susceptible to the new vulnerability.
- Perform risk analysis
- Develop a response plan if warranted
    - If patch available
        - Deploy the patch
        - Update the asset repository
    - If patch not yet available
        - Determine whether interim risk mitigation is required
        - If required, define and apply risk mitigation measures to the target systems; if not, await patch
        - Update the asset repository if required
- Document actions taken
- Assess success/failure of process
- Assess success/failure of process

Proactive Process for Vulnerability Identification and Response

The only difference in this process is that it is proactively initiated as a result of vulnerability assessment scanning. It is identical after the first step.

- Perform vulnerability assessment scanning.
- Receive report of potential vulnerabilities.
- Query the asset repository looking for target systems that are susceptible to the new vulnerability.
- Perform risk analysis
- Develop a response plan if warranted
    - If patch is available
        - Deploy the patch
        - Update the asset repository
    - If patch is not yet available
        - Determine whether interim risk mitigation is required
        - If required, define and apply risk mitigation measures to the target systems; if not, await the patch
        - Update the asset repository if required
- Document actions taken
- Assess success/failure of process

Event Management

Event management provides a process for day-to-day management of the security-related events generated and logged from a variety of sources within the operational environment, including security, network, storage, and host devices. The following processes are required:

- Security logs must be consolidated and maintained. A strategy for storage and maintenance of log files must be defined and implemented.
- Events must be aggregated, normalized, and analyzed regularly to provide a baseline. An event console strategy must be defined and implemented.
- Alerts must be generated and routed to the appropriate individuals when suspicious activity has been detected. In addition, if the event represents an immediate or imminent security threat, then incident management processes must be invoked (see below).

Incident Management

Incident[19] management provides a process framework for responding to security-related threats. Incident management processes are invoked when the analysis of events indicates a violation or imminent threat of violation of computer security policies, acceptable use policies, or standard computer security practices. These include (but are not limited to):

[19] "An incident refers to a computer security problem arising from a threat. Computer security incidents can range from a single virus occurrence to an intruder attacking many networked systems, or such things as unauthorized access to sensitive data and loss of mission-critical data." (NIST)

- Attempts (failed or successful) to gain unauthorized access to a system or its data.
- Unwanted disruption or denial of service.
- The unauthorized use of a system for the processing or storage of data.
- Changes to system hardware, firmware, or software characteristics without the owner's knowledge, instruction, or consent.

Incident management processes include the following:

- A process for analyzing and responding to incidents (see the diagram below).
- A process for contacting appropriate personnel.
- A process for understanding and fulfilling legal requirements (if applicable) with provisions for:
  - Chain of custody
  - Notification of appropriate authorities, including the ability to provide appropriate documentation
  - Recovery
  - Reporting
- A recovery process to bring the organization back to its defined state.
- A reporting process to ensure all interested parties are apprised of incident management activities.

As these descriptions indicate, event and incident management are closely related. Figure 19 shows their relationship graphically.

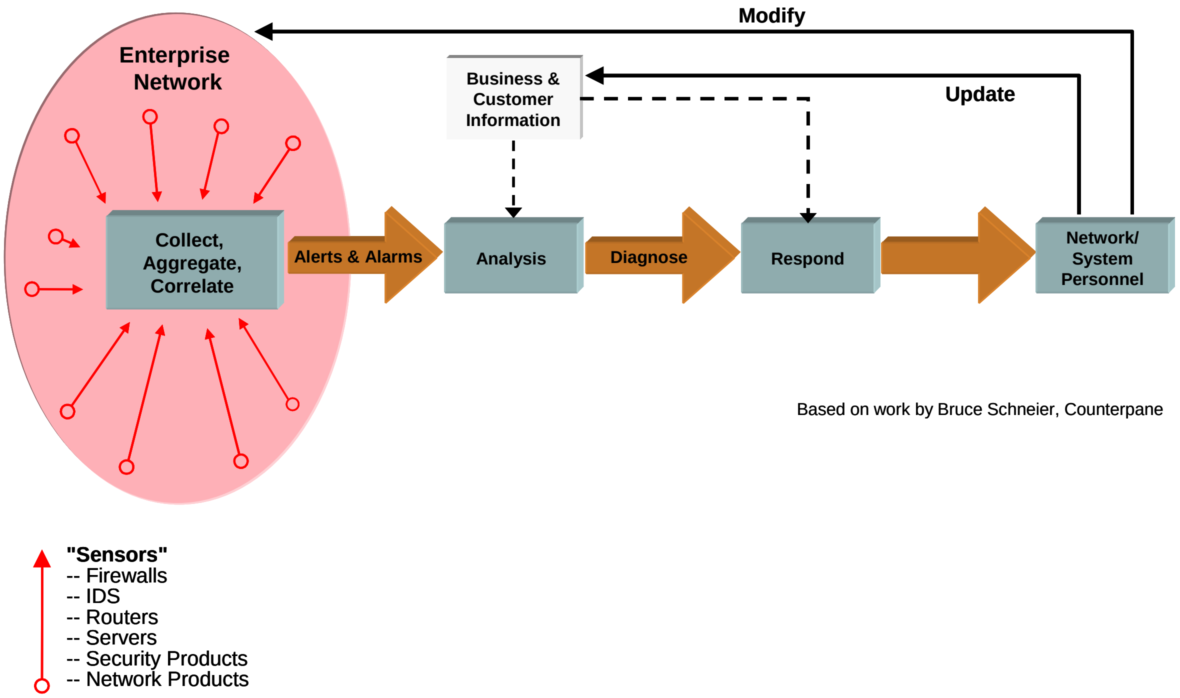

Based on work by Bruce Schneier, Counterpane

Figure 19. Incident Management

This concludes the description of security operations, the third and final component of NAC's ESA. As discussed earlier, security operations encompasses two critical types of processes required to make policy driven security a reality: the processes required to ensure that technology as deployed *conforms to policy* and adequately protects the environment, and the processes required to *enforce policy* within the environment. Although these processes are defined at only a high level, they are equally as important as the other components of ESA. It is in fact these processes that bring policy driven security architecture to life.

Toward Policy Driven Security Architecture

Ideally, a policy-driven security environment is one where policies are articulated at a business level and automatically codified and enforced across the environment. For example, in a healthcare environment the business policy requirement might be expressed simply as “HIPAA[20] compliance”. Our security systems would be able to automatically translate that business policy into the security policies and detailed technical standards required to implement the business policy, and then push them out in electronic form to the various policy decision and enforcement points in the enterprise. In addition, our security systems would have access to the necessary identity and information attributes such that policy decisions could be properly based (e.g., on the characteristics of the requesting user, the requested information, and the environment). Obviously today’s technology does not provide such capabilities, and although there is industry movement in this direction, at current course and speed realization of the full vision may be years away, at best. The policy automation vision, technical model and roadmap discussed below are viewed as important steps in organizing the user and industry actions required to actualize the vision sooner.

Policy Layers and Relationships

Figure 20 briefly defines policy layers and their relationships in the context of the NAC’s ESA conceptual framework for policy-driven security.

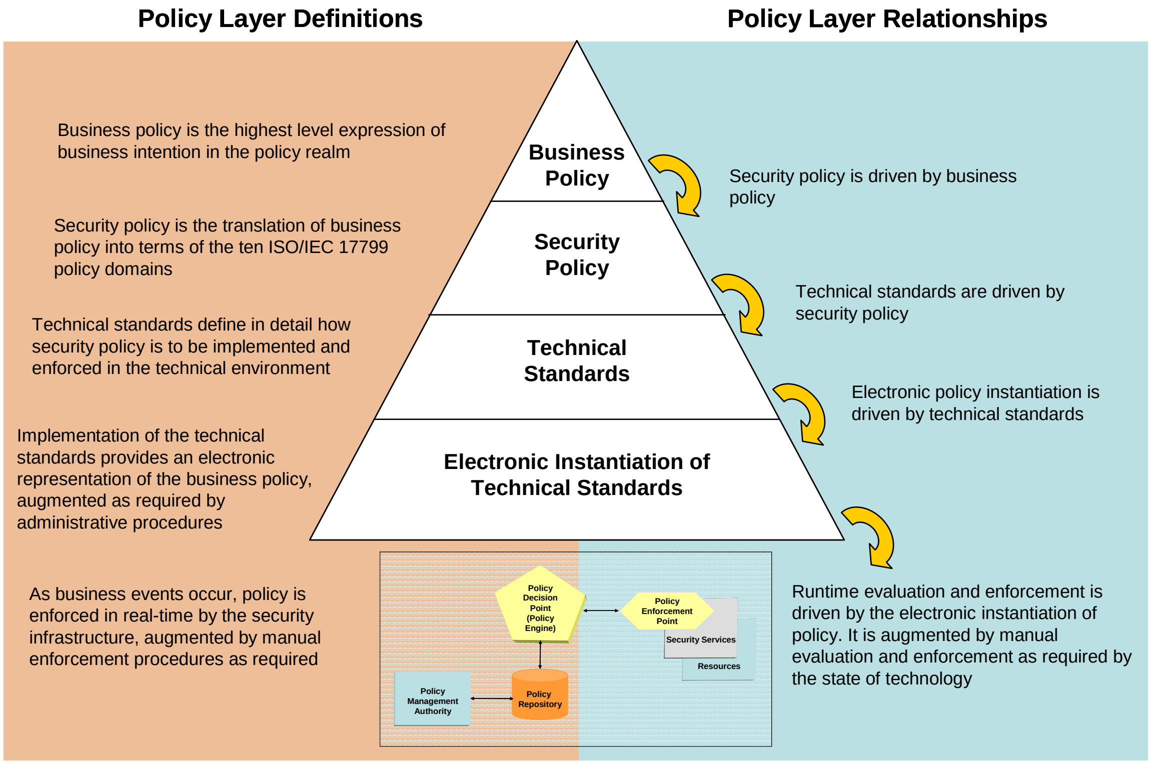

Business policy is the highest level expression of business intention in the policy realm

Security policy is the translation of business policy into terms of the ten ISO/IEC 17799 policy domains

Technical standards define in detail how security policy is to be implemented and enforced in the technical environment

Implementation of the technical standards provides an electronic representation of the business policy, augmented as required by administrative procedures

As business events occur, policy is enforced in real-time by the security infrastructure, augmented by manual enforcement procedures as required

Security policy is driven by business policy

Technical standards are driven by security policy

Electronic policy instantiation is driven by technical standards

Runtime evaluation and enforcement is driven by the electronic instantiation of policy. It is augmented by manual evaluation and enforcement as required by the state of technology

Figure 20. Policy Layers and Relationship

Business policy is the highest level expression of business intention in the security policy realm. The example used is “HIPAA compliance”, which means compliance with the HIPAA standards designed to assure the security of electronic protected health information. One can imagine other examples that might be less industry specific, such as “perimeter access policy” or “software configuration policy”. In the ESA policy model, business policy is then translated to high-level security policies in the relevant ISO/IEC 17799 policy domains. The specific security policies must then be translated into the management and technical standards that define in detail how each policy will be implemented and enforced in the user organization’s technical environment:

- ISO/IEC 17799 Policy Template Section 4 (Organizational Security) is an example of a management policy. For organizations that outsource some of their information processing and support multiple third party access arrangements, these policies may result in the implementation of a number of management standards that define in detail the relevant roles, responsibilities and processes for dealing with service providers and third party access requirements. The management standards themselves are outside the scope of what is being addressed here, keeping in mind that those policy domains such as third party access may spawn technical standards as well.
- ISO/IEC 17799 Policy Template Section 9 (Access Control) and Section 10 (Systems Development and Maintenance) are examples of technical policies. Depending on the scope and diversity of the technical environment, technical policies may translate to a very large number of technical standards[21], defining in detail how each policy is to be implemented and enforced across the technical infrastructure.

Implementation of the technical standards results in an electronic representation of the business policy, augmented as required by administrative procedures. As business events occur, policy is enforced in real-time by the security infrastructure, augmented by manual enforcement procedures as required.

[20] HIPAA is the acronym for the Health Insurance Portability and Accountability Act of 1996, which mandated the establishment of national standards to protect electronic health information.

[21] The term technical standard refers to the standards that implement an organization’s security policies, as identified in an organization-specific policy template. It should not be confused with de jure or de facto industry standards.

Policy Automation Vision

With the above definitions as background, Figure 21 describes the current state and future vision for business policy implementation and enforcement.

As shown on the left of the figure, today's *current state* is that essentially all of the policy and standards definition process is manual, as well as much of the standards implementation process. Implementation is often based on configuration checklists that are associated with each technical standard and each type of hardware or software system where it is to be applied. This may involve manual administration tasks at each system or it may be partially automated based on profiles for each specific system configuration. The amount of manual versus automated configuration definition varies widely from organization to organization.

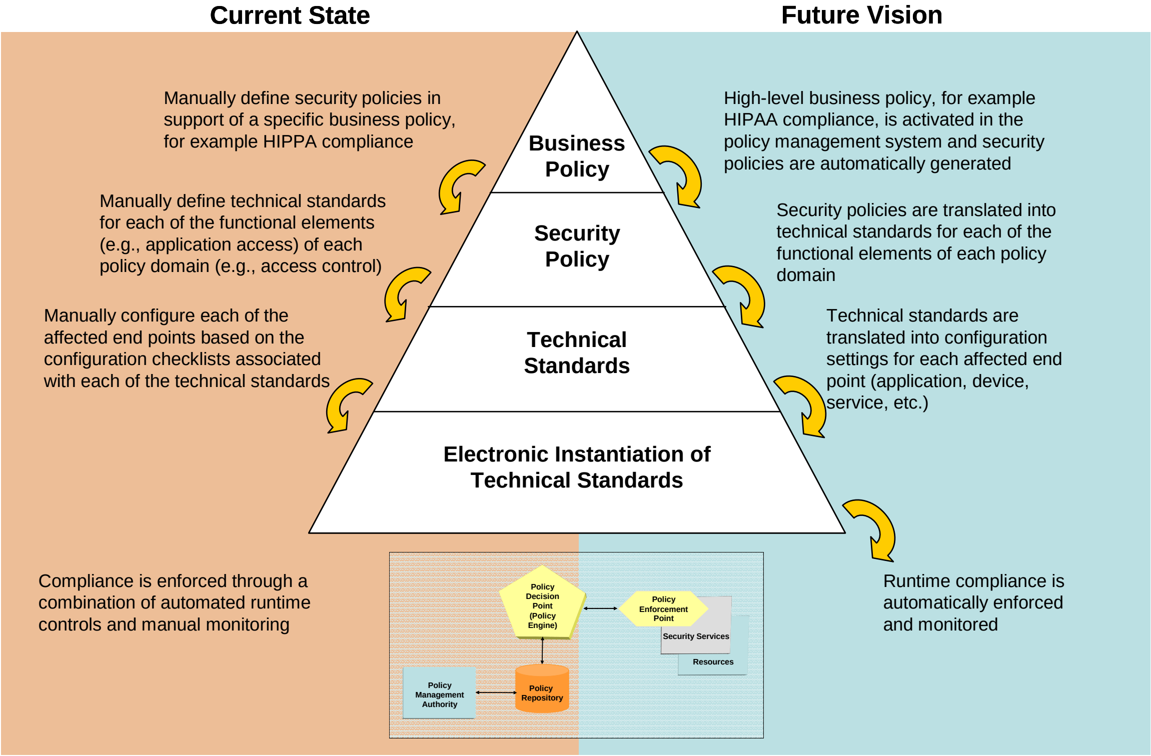
Current State
Future Vision
Manually define security policies in support of a specific business policy, for example HIPPA compliance
Manually define technical standards for each of the functional elements (e.g., application access) of each policy domain (e.g., access control)
Manually configure each of the affected end points based on the configuration checklists associated with each of the technical standards
Compliance is enforced through a combination of automated runtime controls and manual monitoring
Business Policy
Security Policy
Technical Standards
Electronic Instantiation of Technical Standards
High-level business policy, for example HIPAA compliance, is activated in the policy management system and security policies are automatically generated
Security policies are translated into technical standards for each of the functional elements of each policy domain
Technical standards are translated into configuration settings for each affected end point (application, device, service, etc.)
Runtime compliance is automatically enforced and monitored
Policy Decision Point (Policy Engine)
Policy Enforcement Point
Security Services
Resources
Policy Management Authority
Policy Repository

Figure 21. Policy Automation Vision

Once configuration is accomplished, compliance is enforced through a combination of automated runtime controls and manual monitoring. However, there are many cases today in which few if any automated runtime controls are available. This is especially true for standards related to appropriate personal use of equipment and information, which are enforced largely through manual monitoring if at all. Enterprise rights management technology may provide automated runtime controls for appropriate use of information (documents, email, and removable media containing sensitive information) in the future.

As shown on the right of Figure 21, the *future vision* automates all of the policy and standards definition process, as well as the standards implementation process. This is obviously a tall order, and at this point the future arrows on the right of the diagram represent magic.

The next section describes a model for automation of the policy generation and instantiation process, which begins to lay out a technical vision for how the future arrows could potentially be made real. This is followed by a roadmap of user and industry actions that need to occur in order to enable the technical vision.

Policy Automation Model

Figure 22 portrays a high level technical model for automation of the business policy implementation process. At the center is a policy management system with integrated policy mapping, policy translation and technical standards instantiation modules. On the right is a business policy module that provides a generic definition of the business policy to be implemented. On the left is the enterprise-specific policy schema and configuration data required to map the generic business policy definition to the organization's particular technical architecture. The following describes the model in a little more detail, before moving on to an example:

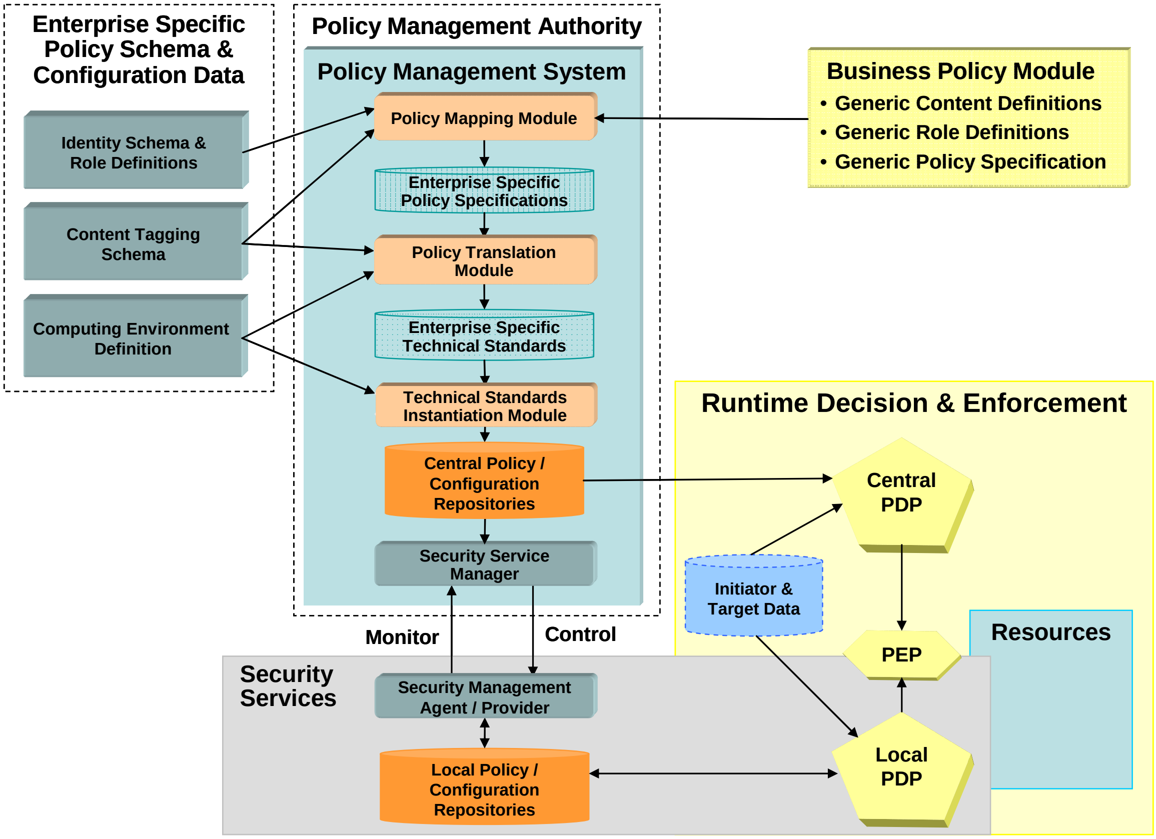
Enterprise Specific Policy Schema & Configuration Data
Identity Schema & Role Definitions
Content Tagging Schema
Computing Environment Definition
Policy Management Authority
Policy Management System
Policy Mapping Module
Enterprise Specific Policy Specifications
Policy Translation Module
Enterprise Specific Technical Standards
Technical Standards Instantiation Module
Central Policy / Configuration Repositories
Security Service Manager
Business Policy Module
• Generic Content Definitions
• Generic Role Definitions
• Generic Policy Specification
Runtime Decision & Enforcement
Central PDP
Initiator & Target Data
PEP
Resources
Monitor
Control
Security Services
Security Management Agent / Provider
Local Policy / Configuration Repositories
Local PDP

Figure 22. Policy Automation Model

- Business Policy Module: It is assumed that there are three major aspects to the generic definition of a business-level policy
  - Generic business content definitions for the particular type of target services / resources affected by the business policy. The automation model example will make this a little clearer.
  - Generic role definitions for the users (or initiators) affected by the business policy.
  - Generic specification of the discrete security policy statements that implement the business policy. It is assumed that these are provided in

the form of a standard policy language that is compliant with the ISO/IEC 17799 policy template.

- Enterprise-Specific Policy Schema & Configuration Data: It is assumed that the following types of schema and configuration data will be required by the policy management system
    - Identity schema and role definitions required to map the generic role definitions to the enterprise specific roles of the particular organization.
    - Content tagging schema required to map the generic content definitions to the enterprise specific definitions for the particular organization.
    - Computing environment definitions (servers, firewalls, directories, etc.) required for policy translation to and technical standards instantiation in the organization-specific environment.
- Policy Management System: The three modules shown represent the three conceptual steps required to map high-level business policy statements to the electronic representation required for runtime policy decision making and enforcement. The security services manager is the management component required to update each of the managed security systems that are affected.
    - Policy Mapping Module: Takes the generic role and content definitions associated with the generic policy specification and maps it to the enterprise specific schema to produce an enterprise specific policy specification. In short, this module maps generic schema to enterprise specific schema.
    - Policy Translation Module: Takes the enterprise specific policy specification statements and translates them based on the enterprise computing environment definition to produce the enterprise-specific technical standards. Note that each policy statement may translate to multiple implementing standards, depending on the number of technical controls, end-point architecture variations and environment-specific variations required. In short, this module converts high-level security policy to detailed technical standards based on the type of device/service.
    - Technical Standards Instantiation Module: Takes the enterprise-specific technical standards and instantiates them in centralized policy and configuration repositories[22] based on the enterprise specific configuration definition. In short, this module instantiates technical standards for each instance of a particular device/service type.
    - Security Service Manager: The security service manager is responsible for updating local policy/configuration repositories for the affected security services. This is accomplished through interactions with the

[22] Although configuration and policy repositories may be distinct, they interact. For example, policy will often regulate what software gets installed and how configuration parameters are defined for a particular device/service type. Policy may also vary for the same device/service type based solely on environment (in which the particular device/service instance is operating at a particular point it time).

security management agents at each of the managed systems[23]. Once the local repositories have been updated, the security services environment is ready for runtime decision making and enforcement.

- Runtime Decision & Enforcement[24]: This portion of the diagram shows the relationship to runtime policy decision making and enforcement components of each ESA security service. Some policy decision points (Central PDPs) may operate directly off of the central policy repositories. It is assumed that these are full function PDPs that are not tightly integrated with the policy enforcement point (PEP) and service implementation. Local PDPs may be less capable, may be tightly integrated with proprietary policy/configuration store and PEP implementations, or may be tightly integrated for performance reasons. As a result, local security management agents may have additional responsibilities, such as mapping policy/configuration updates to proprietary interfaces.

[23] The monitor and control functions represent a standard interface between management systems and managed systems. Monitor is used to determine current state (may be requested or asynchronously reported, depending on implementation). Control is used to update the current state.

[24] For background information, see the Conceptual Framework for Policy-Driven Security and Conceptual Architecture for Policy-Driven Security sections of the ESA document.

Policy Automation Model – HIPAA Example

This section builds on the general automation model with a specific example of business policy implementation. As introduced earlier, the business policy example shown in Figure 23 is HIPAA compliance.

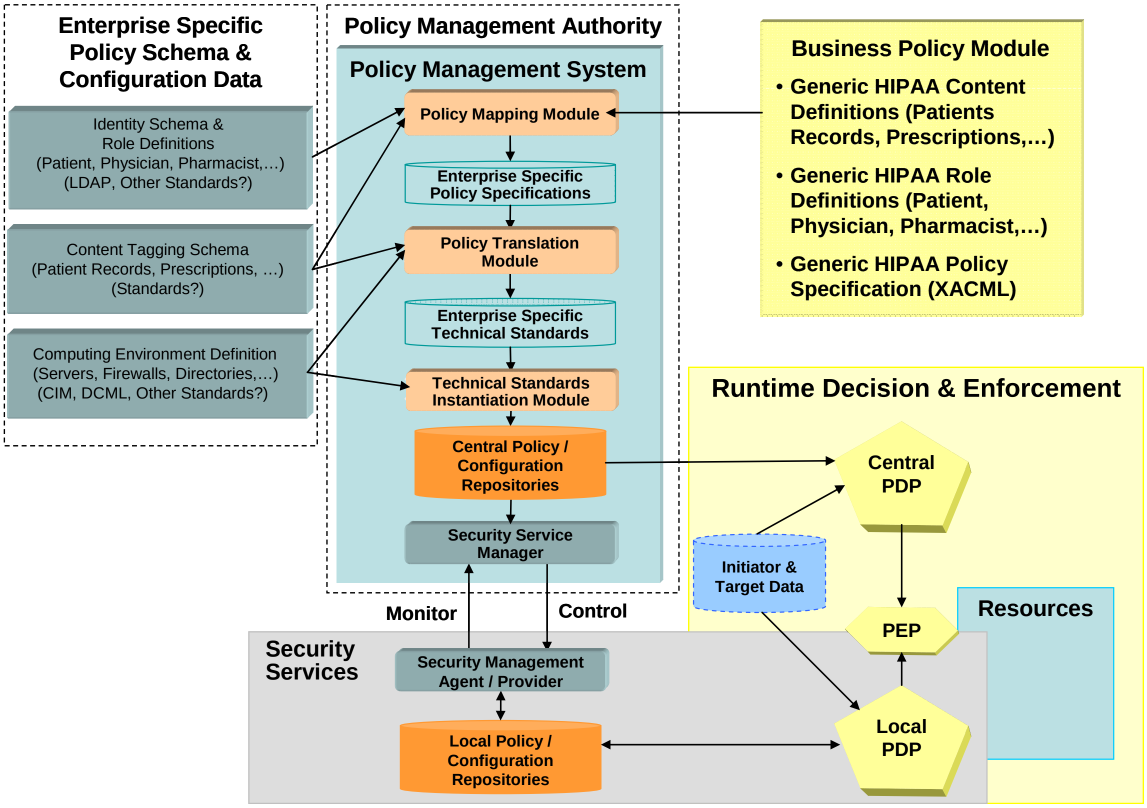

Enterprise Specific Policy Schema & Configuration Data
Identity Schema & Role Definitions (Patient, Physician, Pharmacist,…) (LDAP, Other Standards?)
Content Tagging Schema (Patient Records, Prescriptions, …) (Standards?)
Computing Environment Definition (Servers, Firewalls, Directories,…) (CIM, DCML, Other Standards?)
Policy Management Authority
Policy Management System
Policy Mapping Module
Enterprise Specific Policy Specifications
Policy Translation Module
Enterprise Specific Technical Standards
Technical Standards Instantiation Module
Central Policy / Configuration Repositories
Security Service Manager
Business Policy Module
• Generic HIPAA Content Definitions (Patients Records, Prescriptions,…)
• Generic HIPAA Role Definitions (Patient, Physician, Pharmacist,…)
• Generic HIPAA Policy Specification (XACML)
Runtime Decision & Enforcement
Central PDP
Initiator & Target Data
PEP
Resources
Local PDP
Monitor
Control
Security Services
Security Management Agent / Provider
Local Policy / Configuration Repositories

Figure 23. Policy Automation Model – HIPAA Example

On the right is the HIPAA business policy module, on the left is the enterprise-specific policy schema and configuration data required to map the generic HIPAA policy definition to the organization's particular technical architecture and in the center is the policy management system. The following describes the components of the HIPAA policy automation example in more detail:

- Business Policy Module: It is assumed that the HIPAA module would be provided by or in conjunction with an industry organization that has the HIPAA expertise required to define the generic HIPAA content definitions, role definitions and policy specifications shown. It is also assumed that the generic HIPAA policy specifications are provided in the form of a standard policy language, in this case XACML[25]. Other business policy domains may require policy language standards in addition to those provided by the access control language of XACML.

[25] Extensible Access Control Markup Language (XACML) is an XML-based language, or schema, designed specifically for creating policies and automating their use to control access to disparate devices and applications on a network.

- Identity Schema and Role Definitions: This is the enterprise specific identity schema required to map the generic HIPAA schema to the enterprise specific roles of the particular organization. It is assumed that both sets of schema definitions are provided in a de jure or de facto standard form understood by the policy mapping module.
- Content Tagging schema: This is the enterprise content tagging schema required to map the generic schema for Patient Records, Prescriptions, and so forth to the enterprise specific definitions for the particular organization. It is assumed that content tagging information is provided in a de jure or de facto standard form understood by the policy mapping module.
- Computing Environment Definition: This is the enterprise specific technical computing environment definition (for servers, firewalls, directories, etc.) required for policy translation and technical standards instantiation. It is assumed that this information is provided in a standard form defined by CIM.
- Policy Mapping Module: Takes the generic HIPAA role and content schema definitions associated with the generic policy specification and maps it to the enterprise-specific schema to produce the enterprise-specific HIPAA policy specification. It is assumed that these are in the form of a standard policy language, in this case XACML, and that they are compliant with the ISO/IEC 17799 policy template. Again, other business policy domains may require policy language standards in addition to those provided by the access control language of XACML.
- Policy Translation Module: Takes the enterprise-specific HIPAA policy specification statements and translates them based on the enterprise computing environment definition to produce the enterprise-specific HIPAA technical standards. As discussed earlier, each HIPAA policy statement may translate to multiple implementing technical standards (multiple technical controls for each device/service type).
- Technical Standards Instantiation Module: Takes enterprise-specific HIPAA technical standards and instantiates them in centralized policy/configuration repositories for each instance of a particular device/service type involved in enforcing HIPAA business policy.

Policy Automation Roadmap

With the business policy automation technical model as background, Figure 24 lays out a roadmap of user organization and industry actions required to actualize the technical vision. As the legend indicates, boxes identifying user organization actions briefly describe “Conditions inhibiting automation” on the left and “Conditions supporting automation” on the right. The following describes each of the user organization and industry actions in more detail, starting at the top left:

- Policy: If your organization’s security policy is fragmented, undocumented, and inconsistent without a clear linkage to business-level policy definitions, then there is much that can be done to put yourself in a better position to support policy automation. Start by applying the ESA policy framework - identify the high-level, security-related business principles that will drive your organization’s tailoring of the ISO/IEC 17799 policy template. These high-level, security-related business principles are the business-level policy definitions. Some may be driven by industry specific regulatory requirements such as the HIPAA example. Others may be more organization-centric, such as business-level definitions of “perimeter access policy” or “software configuration policy”. Some may vary based on the specific business unit. The end goal is documented *business policy* and *security policy* consistently translated to *technical standards* for each element of your enterprise-specific technical environment.
- Identifier Semantics Standards: Today there is no industry agreement on how to identify “real world” users, applications, services and resources that are the initiators and targets of requests in NAC’s policy-driven security model. Identifiers vary in syntax and semantics – as exhibited by the use of user-friendly names on one hand, and algorithmically assigned and opaque GUIDs on the other. The following short list describes several problematic aspects of current identifier practices:
    - Representation of an entity’s identifier may take the form of a user ID, UUID, OID, public key, e-mail address, distinguished name, or some other form of identifier (or a combination of the former).
    - Complexities caused by myriad identifier syntaxes are compounded by a lack of consensus around desirable identifier characteristics.
    - Desired characteristics are often mutually exclusive – e.g., visually meaningful vs. obscure to protect privacy; verbally conveyable vs. lengthy to ensure uniqueness; static to support personalization vs. dynamic to avoid profiling; standards-based for interoperability vs. innovative with built in functionalities like check digits.
    - All of this makes it difficult to use identifiers across applications and across organizational boundaries.

  It seems clear that identifiers are highly variable, and flexibility **must** be allowed. But, there is also a need for some level of standardization across application and organizational boundaries. NAC is initiating a work group activity with DMTF participation that that will attempt to reach consensus on requirements and develop a proposal that satisfies those requirements.

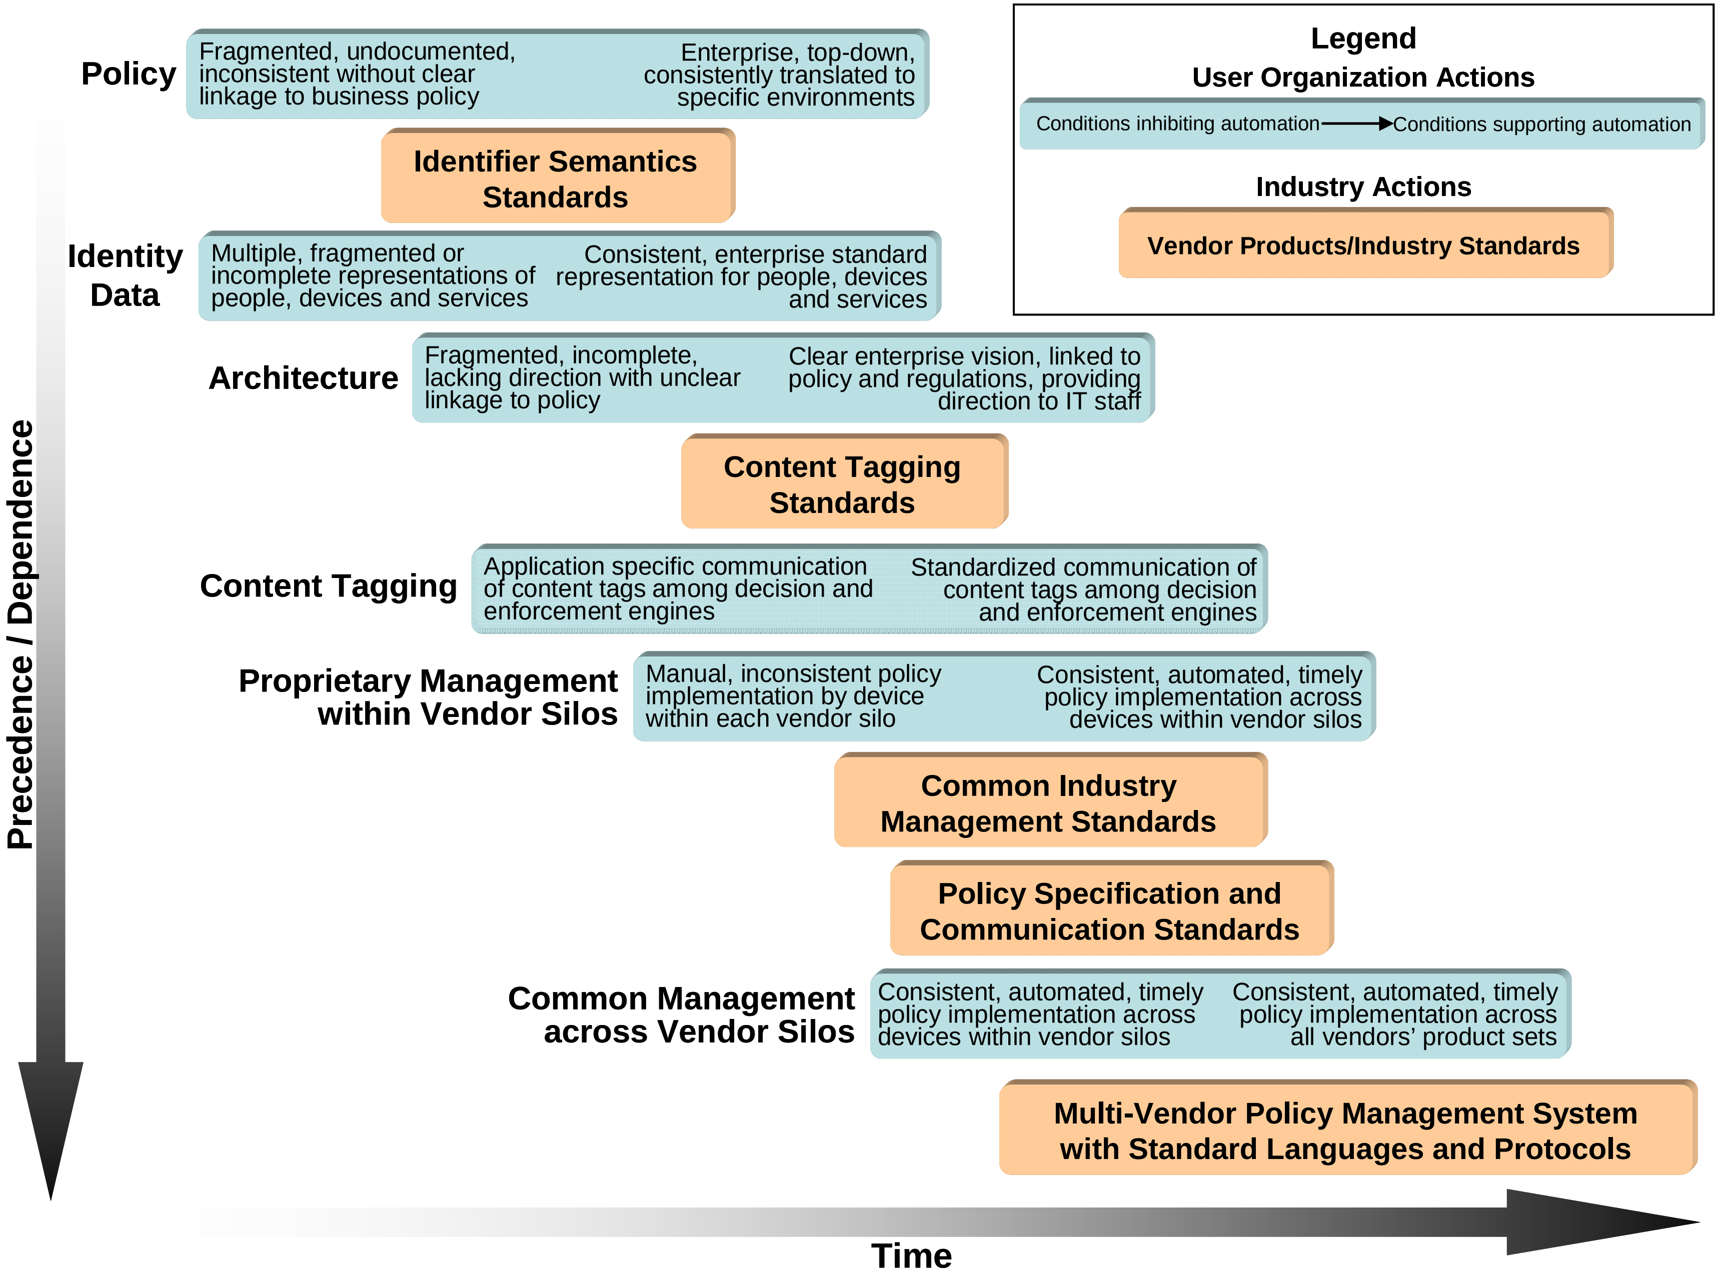
Legend
User Organization Actions
Conditions inhibiting automation → Conditions supporting automation
Industry Actions
Vendor Products/Industry Standards
Policy
Fragmented, undocumented, inconsistent without clear linkage to business policy
Enterprise, top-down, consistently translated to specific environments
Identifier Semantics Standards
Identity Data
Multiple, fragmented or incomplete representations of people, devices and services
Consistent, enterprise standard representation for people, devices and services
Architecture
Fragmented, incomplete, lacking direction with unclear linkage to policy
Clear enterprise vision, linked to policy and regulations, providing direction to IT staff
Content Tagging Standards
Content Tagging
Application specific communication of content tags among decision and enforcement engines
Standardized communication of content tags among decision and enforcement engines
Proprietary Management within Vendor Silos
Manual, inconsistent policy implementation by device within each vendor silo
Consistent, automated, timely policy implementation across devices within vendor silos
Common Industry Management Standards
Policy Specification and Communication Standards
Common Management across Vendor Silos
Consistent, automated, timely policy implementation across devices within vendor silos
Consistent, automated, timely policy implementation across all vendors' product sets
Multi-Vendor Policy Management System with Standard Languages and Protocols
Precedence / Dependence
Time

Figure 24. Policy Automation Roadmap

- Identity Data: If your organization has multiple, fragmented or incomplete schemes for identifying and authenticating users, applications, services and devices, then again there is much that can be done to put yourself in a better position to support policy automation. Digital credential and identity information is the technological foundation for policy-based security. Start by taking every opportunity to move your organization incrementally toward a unified, consistent, enterprise-standard set of identifiers and credentials to be used (possibly with other management data) to authenticate people, devices and services. These credentials should allow trace-back to the identities to which they are assigned. Keep in mind that this step is necessary but may not be sufficient as you begin supporting federated identity management with business affiliates, depending on the extent to which your identifiers, credentials and identity semantics match those of your affiliates. But, having taken this step, your organization will be in a position to more easily move to standard and robust mechanisms for authentication and identity management.

- Architecture: If your organization's security architecture is fragmented and incomplete, perhaps lacking a clear direction and with unclear linkage to policy, then again there is much that can be done to put yourself in a better position. Take advantage of NAC's ESA framework and template for policy-driven security. Start by tailoring the overall enterprise security framework to the needs of your organization, as well as the policy framework, already discussed above. This process will clarify your enterprise security vision, provide a strong linkage to business policy and regulations, and provide direction to IT staff. They can then take advantage of opportunities to tailor specific architectures such as IdM, border protection, and access control to your organization's needs, as well.

- Content Tagging[26] Standards: Currently there is little in the way of standardization activity that one can point to, with the exception of eXtensible Rights Markup Language (XrML), which has its roots in the digital property rights space. To date XrML has not received broad industry support, and the breadth and scope of future support in unclear. In NAC's view, what is needed is a standard way of communicating content tagging information among all of the policy decision and enforcement engines that must support it in order to enable the full set of ESA services. In reality, target data content will be tagged in a variety of ways, and the content tagging information (or tags) will be stored in a variety of ways. Tags may be stored as a field in a database, as a separate database, as a resource attribute field in a directory, as part of a document, or in some other way. A standard way of communicating content tags among policy decision and enforcement points should make this transparent.

- Content Tagging: If your organization's architecture for communicating content tags among policy decision and enforcement points is application specific, then there is work you can do to put yourself in a better position to

[26] Content Tagging is the act of attaching attribute values to digital content via a variety of mechanisms including standard data attributes and metadata. In the context of this document, it is the data (about the target resource) required in order to make and enforce policy decisions to grant or deny requests of the target resource.

support standards-based decision and enforcement engines. This is perhaps best done in the context of a business driver for a standards-based approach, such as a business driver for standards-based access control with business affiliates using SAML. This will allow you to focus an a standard way of communicating content tags among the limited set of policy engines involved in such a project, and will put your organization in a position to more easily move to standardization across a broader set of decision and enforcement engines in the future.

- Proprietary Management within Vendor Silos: If your organization has manual, inconsistent policy implementations by device, then you can begin taking steps to put yourself in a better position to support policy automation. The prerequisite steps are discussed above under Policy and Architecture, but assuming those steps, you may be able to utilize proprietary policy management tools to achieve consistent, automated, timely policy implementation across devices (at least for a specific device vendor). As always, the key is to pick a reputable and forward looking management tool vendor who will support common standards as they become available.

- Common Industry Management Standards: In the context of the policy automation model described earlier, common management standards are of two principle types: standards for describing the computing environment to be managed and standards for management of the environment.

  - Standards for describing an integrated network and computing environment: Currently NAC is aware of several relevant standards activities in this space:

    a) Common Information Model (CIM) from the DMTF - for enterprise and service provider environments

    b) Shared Information/Data model (SID) from the TeleManagement Forum (TMF) - for the telecommunications environment

    c) Standard IETF MIBs such as the Entity MIB (RFC2737) for asset management

    There are other efforts in this space that are currently under development, and should be evaluated as reference documents are released. These include OASIS' Data Center Markup Language (DCML), the Enterprise Grid Alliance (EGA), and GGF's Open Grid Services Architecture (OGSA). Many of these efforts are also being worked through, or aligned with DMTF's CIM.

  - Standards for management of the computing environment

    d) Web-Based Enterprise Management (WBEM) from the DMTF: A set of management and Internet standard technologies developed to unify management of enterprise computing environments. The DMTF has developed a core set of standards that make up WBEM, which includes a data model - the Common Information Model (CIM) standard; an encoding specification for the model - xmlCIM Encoding Specification; and a transport mechanism and set of operations against the model - CIM Operations over HTTP. (Note that the use of XML and HTTP is evolving to a Web services base with the work of the WS-CIM sub team in the DMTF's WBEM Infrastructure and Protocols Working Group.)

e) Web Services Distributed Management (WSDM) from OASIS: The minimum set of management capabilities that should be provided by a resource in a Web services environment. Several management interfaces and discovery are also addressed, as well as how the underlying WS-* foundation is assembled and utilized.

  - Simple Network Management Protocol (SNMP) from the IETF: Currently, there are very few MIBs that provide configuration capabilities, but it is possible to do so. The SNMP protocol is not well-suited to this task and SNMP v3 or later is required in order to support adequate security levels. It is noted that on May 16, 2004, the IETF published an Internet-Draft, Policy Based Management MIB. Interestingly, the document overview is focused on management based on high-level business policies. Also, it should be noted that the DMTF has put emphasis on mapping the standard IETF MIBs to CIM in order to reuse the knowledge that went into the MIB development, and position the MIBs relative to each other and to the other data in CIM, to enable consistent management and policies.

  - Network Configuration Protocol (NetConf) from the IETF: A set of operations for manipulating configuration data sets via a variety of underlying protocols (such as SOAP, BEEP and SSH). Currently, the protocol is model-agnostic, although discussions are underway regarding the formation of a NetConf Data Model Working Group.

- Policy Specification and Communication Standards: In the context of the policy automation model described earlier, there are two requirements: a standard way of specifying policies so that they can be understood and applied across vendor products and security domains, and a standard language for communicating policy decision and enforcement data across vendor products and security domains.

  - Policy specification standards include XACML and the CIM Policy Model. XACML's scope is currently limited to access control, but there are those who believe it would be a good starting point for a more general policy language. On the other hand, a more general language may be less effective than domain-specific languages. In short, the question of whether it could or should be extended requires further study. The CIM Policy Model on the other hand is designed to be independent of any policy language. It appears to be applicable to managing the configuration and behavior of any resource (policy configuration of routers, packet filters, operating systems, storage, etc.), and also addresses authentication and authorization/access control. What are needed are consistent policy and security terminology and a standard that's sufficiently broad in scope to cover all of the ESA services. CIM may provide that coverage, although further work in this area including examination of detailed use cases is needed.

    - Other options for policy specification include languages like PONDER from London’s Imperial College. The coupling of the PONDER syntax with CIM’s semantics has been demonstrated, and further work is underway.
    - Policy communication standards include XACML, SAML and CIM Policy. Again, what is needed is a standard that’s sufficiently broad to encompass communication of policy decision and enforcement data across vendor products and security domains for all ESA services.
    - Lastly, the policy infrastructure itself must be managed. The basic concept of CIM Services could be extended via subclasses to describe and manage the functionality provided by PMAs, PDPs and PEPs. The NAC strongly encourages work in this area, as it will be critical to management of the overall policy infrastructure envisioned by ESA.
- Common Management across Vendor Silos: As the above standards are defined and gain traction, you can begin moving from consistent, automated, timely policy implementation across devices within vendor silos to common policy implementation across all vendors’ devices.
- Multi-Vendor Policy Management System with Standard Languages and Protocols: As the above standards and common management products mature, this enables an enterprise-class, multi-vendor policy management system based on standard languages and protocols.

In summary, NAC’s policy-driven security vision is one in which high-level business policies are automatically translated into the specific security policies and detailed technical standards required to implement the business policy, and then automatically instantiated in a standard form for the various policy decision and enforcement points in the enterprise. An essential corollary is that policy engines must also have access to the necessary identity and management information attributes such that policy decisions can be accurately made (i.e., based on the characteristics of the initiator, the target content, and the environment). Although there is industry movement in this direction, there are standards and technology gaps that must be filled in order to enable the vision across the full set of ESA services and the multiple product and security domains that are involved. Hopefully this business policy automation vision, technical model and roadmap will assist the NAC and alliance partners such as the DMTF in organizing the user and industry actions required to actualize the vision sooner rather than later. The extent to which the full vision can be achieved has yet to be determined, but NAC believes it’s clear that the goal of significantly reducing the manual effort and cost of business policy implementation can be achieved.

Conclusion and Recommendations

As discussed in the executive overview, information systems security has never been more critical or more complex. Complexity is a result of escalating demand for new and improved e-business services in the face of escalating cyber security threats, escalating requirements for corporate governance, and the escalating collapse of traditional protection boundaries. NAC's ESA framework is designed to meet this challenge by simplifying management of this increasingly complex environment. This is accomplished by providing a direct linkage between governance, based on clear and effective business policy, and the security architecture itself.

NAC's security technology architecture focuses on automated policy-driven security, where policy instantiation, decision making and enforcement are built into the architecture. Because the critical standards and implementing products are still immature and incomplete in scope, automated policy-driven security must be considered a work in progress. However, as discussed in detail in the Policy Automation Roadmap section there is a great deal that user organizations can do to better position themselves for future policy automation while at the same time proceeding with an ESA framework that supports partial automation.

Recommendations

This section provides recommendations to user organizations on ways to effectively utilize ESA, both as a common reference for communication and as a starting point for defining their own version of ESA. It also provides recommendations to security infrastructure product vendors and standards organizations for supporting ESA and the policy-driven security architecture vision.

Recommendations to User Organizations

Start using ESA as your common reference architecture for communication on security architecture topics and issues. Use it within your organization and with others in the security space—business partners, vendors, consultants, and industry groups in which you participate.

Start tailoring the ESA framework and template to the needs of your organization:

- Tailor the policy framework to your requirements
    - Define your organization's policy framework, starting with the high-level, security-related business principles that will drive your tailoring of the ISO/IEC 17799 policy template. For more detailed guidance, see the Policy Framework Overview and Policy Automation Roadmap sections of this document. Define detailed management and technical standards, guidelines, and procedures required to implement the policy framework.
    - Implement policy instantiation through processes that minimize manual configuration of decision and enforcement end points to the extent possible. Utilize well defined configuration checklists from NIST or other sources to facilitate centralized configuration definition for each

architectural variant. Implement processes to push proven definitions out to the end points and ensure that the end points are kept in synch.

  - Implement policy enforcement through code procedures that are built into the policy decision and enforcement points, or separate processes where appropriate, recognizing that they will have to be re-engineered as you move to automated policy instantiation and enforcement products.

- Tailor the technology architecture and operations to your requirements, and flesh them out as required to meet the needs of your organization.
- Make sure that you have a quality IdM source and unique-identity strategy in place as a starting point.

Assess business drivers for policy automation products that could further automate both the instantiation and enforcement of policy within a particular context. If the business drivers are there and a reputable standards-based product[27] is available, don't wait—begin incremental implementation so that you can gain hands-on experience with the technology. But, do it in the context of your tailored version of both the ESA framework and the specific ESA architectures such as IdM and border protection that are within the scope of your project. Tailor the architecture to your needs and start building it incrementally around the identified business drivers and products you select for your project. See the Policy Automation Roadmap section for further guidance on how to better position your organization for policy automation

Through your procurement processes encourage ESA vendors to embrace standards-based interoperability and to participate in development and adoption of standards that support the policy automation vision.

Recommendations to Vendors and Standards Organizations

Start utilizing NAC's ESA as a common reference for semantics and terminology around policy-driven security architecture and the enterprise security architecture framework in general. Adopting the terminology used in this document to describe your products and strategies would be valuable to customers and potential customers as they sort through the options offered in the marketplace.

Participate in the development and adoption of standards that support the policy automation vision. For additional detail on the automation vision, models, and roadmap, see the Toward Policy Driven Security Architecture section of this document on page 80. Key standards in the policy-driven security arena include:

- CIM: The Common Information Model is a conceptual information model for describing management that is not bound to a particular implementation.
- CIM Policy Model: A policy language and implementation independent model for conveying general event-condition-action policy semantics.

[27] As suggested in the Policy Automation Roadmap section of the document, proprietary management products are available in some security service and product domains that facilitate automation across a particular vendor's product set.

- DCML: The Data Center Markup Language is an emerging standard for describing the computing environment to be managed.
- ISO/IEC 17799:2000: An international standard for information security management that is gaining traction in the enterprise security space. NAC selected it as an integral part of the ESA policy framework based on industry recommendations received at the NAC 2004 Spring Conference.
- LDAP: The standards-based means for accessing identity authentication and authorization data and related policy data that are stored in an X.500 directory.
- SAML: The Security Assertion Markup Language provides the standards-based means for communication of identity, attributes, and authorization decisions related to initiators and targets.
- SNMP: The Simple Network Management Protocol is the most pervasive management standard today, broadly supported and used for network configuration management.
- WBEM: Web-Based Enterprise Management is a set of management and Internet standard technologies developed to unify management of enterprise computing environments.
- WS-Policy: An emerging standard for specifying initiator/target interaction policies in a Web services environment (for example, initiator must authenticate using X.509 certificate or using SAML).
- X.509: X.509 is the fundamental public key infrastructure-based technology critical for establishing identities and secure, trusted communications between the components. In many cases, X.509 certificates may be used in place of SAML assertions to provide initiator identity.
- XACML: The Extensible Authorization Control Markup Language provides the standards-based means to specify and communicate access control policy. XACML is used for communication of policy between the policy repository and the PDP.

Product vendors should consider the opportunities afforded by policy-based security architecture in general, and by the automated policy instantiation and enforcement vision in particular. The products that will thrive are those based on open standards and a common vision, with product differentiation based on interoperability features, standards-based functionality, performance, and reliability rather than proprietary management features that attempt to lock users into a particular vendor silo.

Appendix A. Glossary of Terms

Many of the definitions in this glossary are taken from *NIST Special Publication 800-33, Underlying Technical Models for Information Technology Security, December 2001*. A few are the NAC's own definitions adapted from other sources.

Access control

Enable authorized use of a resource while preventing unauthorized use or use in an unauthorized manner.

Accountability

The security objective that generates the requirement for actions of an entity to be traced uniquely to that entity. This supports non-repudiation, deterrence, fault isolation, intrusion detection and prevention, and after-action recovery and legal action.

Assurance

Grounds for confidence that the other four security objectives (integrity, availability, confidentiality, and accountability) have been adequately met by a specific implementation. "Adequately met" includes (1) functionality that performs correctly, (2) sufficient protection against unintentional errors (by users or software), and (3) sufficient resistance to intentional penetration or by-pass.

Authentication

Verifying the identity of a user, process, or device, often as a prerequisite to allowing access to resources in a system.

Authorization

The granting or denying of access rights to a user, program, or process.

Availability

The security objective that generates the requirement for protection against intentional or accidental attempts to (1) perform unauthorized deletion of data or (2) otherwise cause a denial of service or data.

Confidentiality

The security objective that generates the requirement for protection from intentional or accidental attempts to perform unauthorized data reads. Confidentiality covers data in storage, during processing, and while in transit.

Data integrity

The property that data has not been altered in an unauthorized manner. Data integrity covers data in storage, during processing, and while in transit.

Data origin authentication

The verification that the source of data received is as claimed.

Denial of service

The prevention of authorized access to resources or the delaying of time-critical operations.

Domain

See security domain.

Entity

Either a subject (an active element that operates on information or the system state) or an object (a passive element that contains or receives information).

Guideline

A guideline is an enterprise-wide recommended course of action. While not mandatory, it is highly encouraged that guidelines be reviewed for applicability to particular environments, and implemented as appropriate for the business environment. Guidelines support the policy and the standards.

Integrity

The security objective that generates the requirement for protection against either intentional or accidental attempts to violate data integrity (the property that data has not been altered in an unauthorized manner) or system integrity (the quality that a system has when it performs its intended function in an unimpaired manner, free from unauthorized manipulation).

Identity

Information that is unique within a security domain and is recognized as denoting a particular entity within that domain.

Identity-based security policy

A security policy based on the identities and/or attributes of the object (system resource) being accessed and of the subject (user, group of users, process, or device) requesting access.

Incident

A violation or imminent threat of violation of computer security policies, acceptable use policies, or standard computer security practices. These include (but are not limited to):

- Attempts (failed or successful) to gain unauthorized access to a system or its data
- Unwanted disruption or denial of service
- The unauthorized use of a system for the processing or storage of data
- Changes to system hardware, firmware, or software characteristics without the owner's knowledge, instruction, or consent

IT-related risk

The net mission/business impact (probability of occurrence combined with impact) from a particular threat source exploiting, or triggering, a particular information technology vulnerability. IT related-risks arise from legal liability or mission/business loss due to:

- Unauthorized (malicious, non-malicious, or accidental) disclosure, modification, or destruction of information.
- Non-malicious errors and omissions.
- IT disruptions due to natural or man-made disasters.
- Failure to exercise due care and diligence in the implementation and operation of the IT.

IT Security Architecture

A description of security principles and an overall approach for complying with the principles that drive the system design; i.e., guidelines on the placement and implementation of specific security services within various distributed computing environments.

Policy

A broad statement authorizing a course of action to enforce the organization’s guiding principles for a particular control domain. Policies are interpreted and supported by standards, guidelines, and procedures. Policies are intended to be long-term and guide the development of rules to address specific situations.

Principles

In this document, the agreed-upon set of security principles that govern the use and management of technology across an organization. They are derived from a combination of (1) basic assumptions and beliefs that reflect the organization’s mission, values, and experience; and (2) business, legal, and technical principles that drive the enterprise.

Procedure

A procedure provides instructions describing how to achieve a policy or standard. A procedure establishes and defines the process whereby a business unit complies with the policies or standards of the enterprise.

Risk

Within this document, the term is synonymous with “IT-related risk.”

Risk analysis

The process of identifying the risks to system security and determining the probability of occurrence, the resulting impact, and the additional safeguards that mitigate this impact. It is part of risk management and synonymous with risk assessment.

Risk assessment

See risk analysis

Risk management

The total process of identifying, controlling, and mitigating IT- related risks. It includes risk analysis; cost-benefit analysis; and the selection, implementation, test, and security evaluation of safeguards. This overall system security review considers both effectiveness and efficiency, including impact on the mission/business and constraints due to policy, regulations, and laws.

Rule-based security policy

A security policy based on global rules imposed for all subjects. These rules usually rely on a comparison of the sensitivity of the objects being accessed and the possession of corresponding attributes by the subjects requesting access.

Security

Security is a system property. Security is much more than a set of functions and mechanisms. IT security is a system characteristic as well as a set of mechanisms that span the system both logically and physically.

Security domain

A set of subjects, their information objects, and a common security policy.

Security goal

The IT security goal is to enable an organization to meet all mission/business objectives by implementing systems with due care consideration of IT-related risks to the organization, its partners, and its customers.

Security policy

The statement of required protection of the information objects.

Security objectives

The five security objectives are availability, integrity, confidentiality, accountability, and assurance.

Standard

A standard is an enterprise-wide, mandatory directive that specifies a particular course of action. Standards support the policy and outline a minimum baseline for policy compliance.

System integrity

The quality that a system has when it performs its intended function in an unimpaired manner, free from unauthorized manipulation of the system, whether intentional or accidental.

Threat

The potential for a “threat source” (defined below) to exploit (intentional) or trigger (accidental) a specific vulnerability.

Threat source

Either (1) intent and method targeted at the intentional exploitation of a vulnerability or (2) the situation and method that may accidentally trigger a vulnerability.

Threat analysis

The examination of threat sources against system vulnerabilities to determine the threats for a particular system in a particular operational environment.

Vulnerability

A weakness in system security procedures, design, implementation, internal controls, etc., that could be accidentally triggered or intentionally exploited and could result in a violation of the system’s security policy.

Appendix B. Glossary of Resources

Security Governance Resources and Tools

Policy Development Tools

The ISO/IEC 17799:2000, Code of Practice for Information Security Management, is an international standard that is gaining traction in the enterprise security space. The NAC selected it as its policy template based on industry recommendations at the NAC 2004 Spring Conference. It is believed to have broad applicability across the many different organizational types represented in NAC.

Extracts from BS EN ISO 17799:2000 are reproduced as part of Enterprise Security Architecture: A Framework and Template for Policy-Driven Security document with the permission of BSI under license number 2004AT0134. Hard copies of other British Standards are available from BSI Customer Services, 389 Chiswick High Road, London W4 4AL, United Kingdom. [Tel + 44 (0)20 8996 9001]. E-mail : cservices@bsi-global.com.

In addition to the ISO 17799:2000 standard itself, other resources such as the ISO17799 Toolkit (http://www.iso17799-made-easy.com) are also available. The toolkit contains:

- Both parts of the ISO17799 standard itself[28]
- A full set of ISO17799-compliant information security policies
- A management presentation on ISO 17799 / BS7799 in PowerPoint format
- A disaster recovery planning kit (related to ISO 17799 section 11)
- A road map for certification
- An audit kit (checklists, etc.) for a modern network system (section 12)
- A comprehensive glossary of information security and computer terms
- A business impact analysis questionnaire

There are other standards and tool kits available outside the ISO model. The most notable is the (NIST) 800-26 model. This is a free toolkit but has been identified by some members of the NAC as more complex and more cumbersome to apply. For more information, visit http://csrc.nist.gov/.

For other policy examples and templates, visit the SANS Security Policy Project Web site at http://www.sans.org/resources/policies/. This is a consensus research project of the SANS community whose ultimate goal is to offer rapid development and implementation of information security policies. However, this is an incomplete policy template set.

[28] The second part of the standard is the specification for information security management systems, also known as BS 7799-2:2002.

Metrics

The NIST Special Publication 800-55, Security Metrics Guide for Information Technology Systems, “provides an approach to help management decide where to invest in additional security protection resources or identify and evaluate nonproductive controls. It explains the metric development and implementation process and how it can also be used to adequately justify security control investments. The results of an effective metric program can provide useful data for directing the allocation of information security resources and should simplify the preparation of performance-related reports.”

NIST References for ESA Implementation

The following is a list of recommended “Things to Use Now” from the NIST presentation at the NAC 2004 Spring Conference:

- Tele-Commuting Guidance (***): SP 800-46 Security for Telecommuting and Broadband Communications, September 2002
- Patch Management Guidance (***): SP 800-40 Procedures for Handling Security Patches, September 2002
- Incident Handling (****): SP 800-61, Computer Security Incident Handling Guide, January 2004
- Network Security Testing Guidance (***): SP 800-42 Guideline on Network Security Testing, October 2003
- Web Server Security Guidance (***): SP 800-44 Guidelines on Securing Public Web Servers, September 2002
- E-Mail Security Guidance (***): SP 800-45 Guidelines on Electronic Mail Security, September 2002
- W2K Administration Guidance (****)+: SP 800-43 Systems Administration Guidance for Windows 2000 Professional, November 2002
- Self-Assessment Tool (***): Automated Security Self-Evaluation Tool (ASSET). The purpose of ASSET is to automate the completion of the questionnaire contained in NIST Special Publication SP 800-26 Security Self-Assessment Guide for Information Technology Systems, November 2001

Other more recent NIST publications:

- Draft NIST Special Publication 800-68, Guidance for Securing Microsoft Windows XP Systems for IT Professionals: A NIST Security Configuration Checklist
- Draft NIST Special Publication 800-70, The NIST Security Configuration Checklists Program
- Draft NIST Special Publication 800-72, Guidelines on PDA Forensics
- Draft Special Publication 800-66: An Introductory Resource Guide for Implementing the Health Insurance Portability and Accountability Act (HIPAA) Security Rule

For up to date information, go to NIST Computer Security Special Publications.