The HP Security Handbook

Protecting Your Business

Governance and Compliance
Proactive Security Management
Identity Management
Trusted Infrastructure
Innovation in Information Security

The HP Security Handbook

Vice President, HP Security Office
Tony Redmond

HP Security Handbook Lead
Jan De Clercq

Layout, Illustrations, Graphics and Production Lead
Killian McHugh

Primary Content Editors and Strategists
Governance and Compliance: Stuart Hotchkiss (Lead)
Proactive Security Management: Keith Millar (Lead)
Identity Management: Jan De Clercq (Lead), Mark Crosbie
Trusted Infrastructure: Boris Balacheff (Lead), Iver Band, Archie Reed, Mark Schiller, Bill Wear
Innovation in Security: Simon Shiu (Lead), Joe Pato

Additional Content Contributors

Governance and Compliance:
Lois Boliek, John Carchide, Frederic Gittler, David Graves, Jim Hoover, Cheryl Jackson, Paul Jeffries, Bill Kowaleski, Tari Schreider, Saida Wulteputte, Mike Yearworth

Proactive Security Management:
Hayden Brown, John Carchide, Tracy DeDore, Paul Jeffries, Montserrat Mane, Jim O'Shea, Christopher Peltz, Sarah Porten, Yann Vermast, Brian Volkoff, Doug Young

Identity Management:
Sai Allavarpu, Jean-Michel Argenville, Carolyn Bosco, Pete Bramhall, Christian Fischer, Ronald Luman, Marco Casassa Mont, Robert Neal-Joslin, Jason Rouault, Scott Swist, Ibrahim Wael, Manny Novoa

Trusted Infrastructure:
Enrico Albertin, Shivaun Albright, Sunil Amanna, Mike Balma, Ron Carelli, Lynne Christofanelli, Paul Congdon, Joanne Eames, Janusz Gebusia, Gary Lefkowitz, Shab Madina, Sunil Marolia, John Rhoton, Steve Scott, Rick Supplee, Tom Welsh, Chris Whitener

Innovation in Security:
Adrian Baldwin, Richard Brown, Chris Dalton, Bill Horne, Ed McDonnell, David Pym, Martin Sadler, Steve Simske, Richard Smith

The HP Security Handbook is available online at www.hp.com/go/security/securityhandbook.
For feedback, please e-mail hpsecurityhandbookfeedback@hp.com.
For additional printed copies, please see your HP Sales Representative.

About HP

HP is a technology company that operates in more than 170 countries around the world. We explore how technology and services can help people and companies address their problems and challenges, and realize their possibilities, aspirations, and dreams. We apply new thinking and ideas to create more simple, valuable, and trusted experiences with technology, continuously improving the way our customers live and work.

No other company offers as complete a technology product portfolio as HP. We provide infrastructure and business offerings that span from handheld devices to some of the world's most powerful supercomputer installations. We offer consumers a wide range of products and services from digital photography to digital entertainment and from computing to home printing. This comprehensive portfolio helps us match the right products, services, and solutions to our customers' specific needs.

HP focuses on simplifying technology experiences for all of its customers - from individual consumers to the largest businesses. With a portfolio that spans printing, personal computing, software, services and IT infrastructure, HP is among the world's largest IT companies, with revenue totaling $104.3 billion for the four fiscal quarters ended Oct. 31, 2007. More information about HP is available at **www.hp.com**.

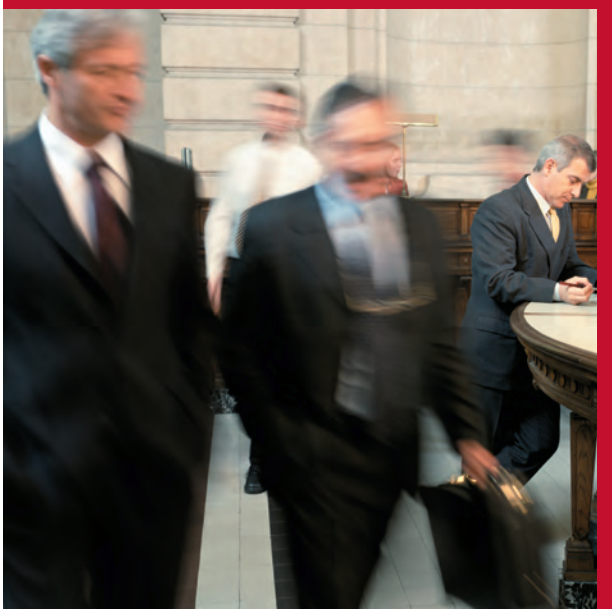

Fast facts

• HP was incorporated in 1939.
• Corporate headquarters are in Palo Alto, California.
• Mark Hurd is president and CEO.
• HP is a US Fortune 14 and a Global Fortune 41 company, with revenue totaling $104.3 billion for the fiscal year ended Oct. 31, 2007.
• HP has 150,000 employees doing business in more than 170 countries around the world.

Technology leadership

HP's three business groups drive industry leadership in core technology areas:

• The Personal Systems Group: Business and consumer PCs, mobile computing devices, and workstations
• The Imaging and Printing Group: Inkjet, LaserJet and commercial printing, printing supplies, digital photography, and entertainment
• The Technology Solutions Group: Business products including storage and servers, managed services, and software

Contribution

HP strives to be an economic, intellectual, and social asset to each country and community in which we do business. Key areas of contribution are electronic waste, raising standards in our global supply chain and increasing access to information technology.

Growth

HP is focused on three technology shifts that have the power to transform our customers' lives and businesses:

• Next-generation data center
• Always on, always connected mobile computing
• Ubiquitous printing and imaging

For more information, visit **www.hp.com**.

About HP's Security Practice

HP takes a holistic approach to security that includes the people, process and technology to ensure the effectiveness of the security solution. HP Services assists in defining a security strategy specifically tailored to the customer's environment and business processes. As a leader in IT management, and more specifically IT service management, HP Services brings tremendous breadth and depth of management expertise to every consulting engagement. Our expert security staff includes Certified Information Systems Security Professionals (CISSPs) and certified Sysadmin, Audit, Network, Security (SANS) individuals who bring extensive experience in multi-vendor platforms including HP-UX, IBM AIX, Sun Solaris, OpenVMS, Microsoft Windows, and Linux. As a member of the Information Technology Information Sharing and Analysis Center (IT-ISAC), HP's security services team stays abreast of the latest information on cyber security issues and utilizes proven best practices and methodologies such as BS 7799/ISO 17799.

Table of contents

Chapter One: Governance and Compliance

Chapter Two: Proactive Security Management

Chapter Three: Identity Management

Chapter Four: Trusted Infrastructure

Chapter Five: Innovation in Information Security

"The HP Security Framework covers the range of security, governance and risk management subjects required of a truly professional security programme. Security is not just a technology issue, neither is it just a single-point problem. Only by stepping back and seeing the whole risk picture can good security be made to work, and I applaud the authors of the Security Handbook in getting this message across."
-Dr Paul Dorey, CISO, BP plc, and Chairman of the Institute of Information Security Professionals (IISP)

Introduction Note

Tony Redmond

Apart from being the world's largest IT company, in many other ways, HP is a unique IT company. No other company develops the same breadth and depth of technology across all market segments - from consumer to enterprise, from small and midsize businesses to the public sector - spanning so many types of computing devices, protocols, standards, and applications. Security is and will remain a prime focus for HP across our complete portfolio because customers expect that everything that they buy from HP is secure.

The HP Security Handbook provides a view of all the different threads of security in which HP works. We plan to update the content regularly; the handbook is an evolving document that tracks new developments, adds new information as it becomes available, and presents industry standards and initiatives important to security as they mature. Much of the content focuses on the three pillars of HP's security strategy - identity management, proactive security management, and trusted infrastructures. These are all "big plays", places where HP believes that we can make a real difference in the way that people use technology.

I cannot think of a bigger challenge than building truly trustworthy infrastructures composed of new hardware architectures, new operating systems, and new applications. Federated identity management will help to liberate users from the tyranny and insecurity of multiple user name and password pairs. And proactive security management is HP's way of declaring to the security industry that it's time to stop reacting to threats and begin building intelligence in servers and other network components to better resist unauthorized intrusions. Because HP is such a large company, we have a special responsibility and role within the IT industry to help chart the future, and that's what the HP security strategy sets out to do. Big plays don't happen overnight - but these plays form the core of our strategy because they are worthwhile and will make a difference.

In addition to describing HP's security strategy, this handbook illustrates the broad sweep of security activity across the range of the company's offerings, from services to the fundamental security features incorporated in HP operating systems. It also describes the work of the Trusted Systems Laboratory and how HP researchers look at future security challenges. Commentators often say that the only guarantee regarding technology is change; this is especially true for security technology. HP's investment in research has already provided great benefits, and we expect this trend to continue.

Tony Redmond
Vice President, HP Security Office

Introduction

Information security is a fundamental necessity and enabler for modern business. Because information technology infrastructures provide the ability for enterprises to automate, adapt, and accelerate their business strategies, information security is essential for safeguarding business continuity. Whether enabling secure sharing and collaboration with partners, preventing or detecting insider attacks, or defending against indiscriminant vandalism by unseen and random network attackers - information security is a key element of any IT infrastructure.

Security, however, is not a simple commodity that can be ordered by weight and bolted on to an IT infrastructure. Security considerations should permeate every aspect of IT - from the design of applications and infrastructure to the mechanisms for managing their deployment; from discrete components that protect specific functions to the design of business objectives and the governance of corporate policy; from the management of technology to the management of people.

Measuring security is also difficult - how safe are we at any point? Unlike processor speed or storage capacity, we do not measure security in simple units - except after an incident when we can objectively demonstrate that the deployed security mechanisms are inadequate. As a result, enterprise security is traditionally mired in a cycle of reactive crises.

The Security Landscape

Enterprises face a rapidly changing environment that demands a proactive stance for information security. Key factors driving this change include:

- Unrelenting presence of security incidents through out the industry
- Ever-increasing sophistication of attack
- Government regulation
- Changes in IT infrastructure to accommodate changing business

Continuing presence of security incidents

High-profile security breaches have made network security one of the most important concerns for corporate and government networks. In the recent past, the rate of security incidents grew at a tremendous rate. More recently, the rate of attacks has leveled off, but many attacks now target specific victims or resources rather than the indiscriminate attacks prevalent earlier in the decade. As reported in the 2007 edition of CSI Computer Crime and Security Survey, the average reported annual loss from security incidents doubled from the previous year.

Increasing sophistication of attacks

The emergence of targeted attacks is coupled with an increase in sophistication of attack. Not only are specific victims selected for attack, but unrelated organizations or individuals are also selected for attack to serve as staging points for stealthy attacks. An underground economy has emerged for access to compromised systems for direct exploitation or for use to stage subsequent attacks.

Government response through regulation

Governments have not ignored the increasing threat to commerce. Many governmental entities have enacted or are preparing legislation to require business attention to information security issues.

Figure i-1
Sample regulations affecting security

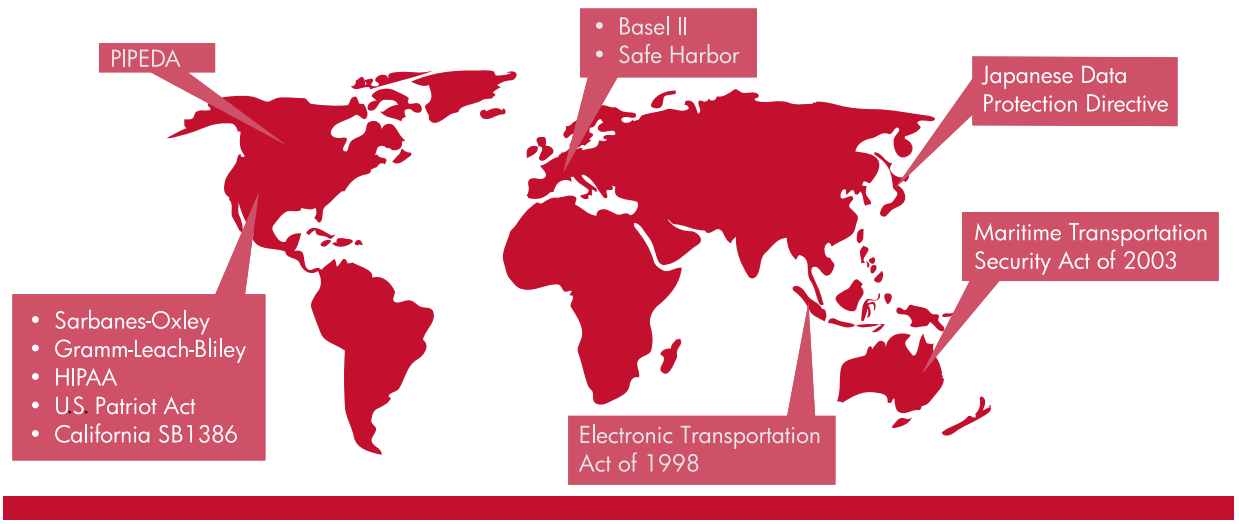

Regulatory mandates such as the Sarbanes-Oxley Act of 2002, the California Database Protection Act of 2001, the Gramm-Leach-Bliley Act (GLB), the Health Insurance Portability and Accountability Act (HIPAA), and the Basel II Accord are an additional catalyst for applying due diligence in the security decision and implementation process. These laws impose strict requirements on enterprises to establish, identify, document, test, and monitor necessary internal control processes. Because information technology supports most, if not all, of these processes, these laws significantly affect companies' security strategies. These new regulations drive security designers and architects to impose and maintain the proper security controls throughout their enterprise.

Changing business objectives and streamlining processes

The need for business agility is driving the development of proactive security capabilities. Ad-hoc security implementations often interlock the various components of a business application, which limits the overall ability to adapt, increases the cost to operate, and often leads to diminishing protection through an application's lifetime. Enabling rapid flexibility requires an overall process for managing and evolving an organization's IT security.

HP's Security Framework

Delivering a safer enterprise IT environment aligned to defined levels of security and risk requires a framework for rapid and effective response to threats and corporate business objective changes. HP's security framework, shown in Figure i-2, enables a holistic way to proactively define and deliver security across the enterprise.

The key areas represented in this model include the three areas in which HP is investing to create innovation and differentiation: identity management, proactive security management, and trusted infrastructure. The fourth area, governance, includes the supporting services and tools that HP delivers to ensure that IT security solutions meet business objectives.

Business Context

The top level of the security framework consists of the key drivers, including business objectives, operational risk, and regulatory and legal compliance. Businesses and organizations have a set of major objectives or missions that drive their existence. In addition, they must manage operational risk and meet regulatory and legal compliance. All of these factors have direct security implications that drive the overall security strategy of a business or organization.

Figure i-2
HP's security framework

From the security perspective, examples of threats that directly affect the highest levels of a company or an organization include:

- Theft of intellectual property or digital assets
- Disruption of critical services or infrastructure that leads to lost revenues, contractual breaches, or regulatory violations
- Public disclosure of sensitive information, which negatively impacts brand identity or competitive advantage

Governance and Compliance

Governance refers to the controls and policies that translate high-level business objectives, operational risks, and regulatory needs into the directives, objectives, and policies that drive security mechanisms. Governance is a strategic component of every technology optimization initiative. It includes business logic, business procedures, managerial processes, and operational processes that are all supported by specific, lower-level policies for IT operations and security.




Proactive Security Management

Proactive security management focuses on managing security functions in support of business and organizational goals and processes. The fundamental goal of this area is to ensure that protection mechanisms operate appropriately during setup, operation, and decommissioning of various IT services. Proactive security management:

- Manages the protection of data, applications, systems, and networks, both proactively and reactively
- Supports changing business and organizational models and responds to a changing threat environment
- Maintains the level of security and operational risk defined by a company or organization

Identity Management

Identity management is the ability to identify every user, application, or device across an organization or business. It provides flexible authentication, access control, and auditing while respecting privacy and regulatory controls.

Delivered via a set of processes and tools for creating, maintaining, and terminating a digital identity, identity management allows administrators to manage large populations of users, applications, and systems quickly and easily. The tools permit selective assignment of roles and privileges, which facilitates compliance with regulatory controls and contributes to privacy-sensitive access controls.

Trusted Infrastructure

Trusted infrastructures are composed of hardware platforms, together with their operating environments and applications, which behave in an expected and predictable way for their intended purpose. Trusted infrastructures must support the IT applications underlying the most critical business processes. When IT infrastructure technologies fail to keep pace with emerging threats, we no longer trust them to sustain the applications we depend on in both business and society.

A trusted infrastructure reliably manages and controls access to information assets while delivering the power required for critical business processes. It helps implement appropriate technologies to secure the end-to-end IT infrastructure of a company or organization, worldwide - including data centers, networks, productivity tools, end-user desktops, and wireless devices.

The need for a trusted IT infrastructure flows from our increasing reliance on IT systems to do everything from running our business to running our society's utilities. Just as our dependence on IT permeates all aspects of society, security capabilities must permeate all aspects of IT infrastructure. Security must be built in, not bolted on, at the platform level, at the network level, and in the very processes used for developing systems.

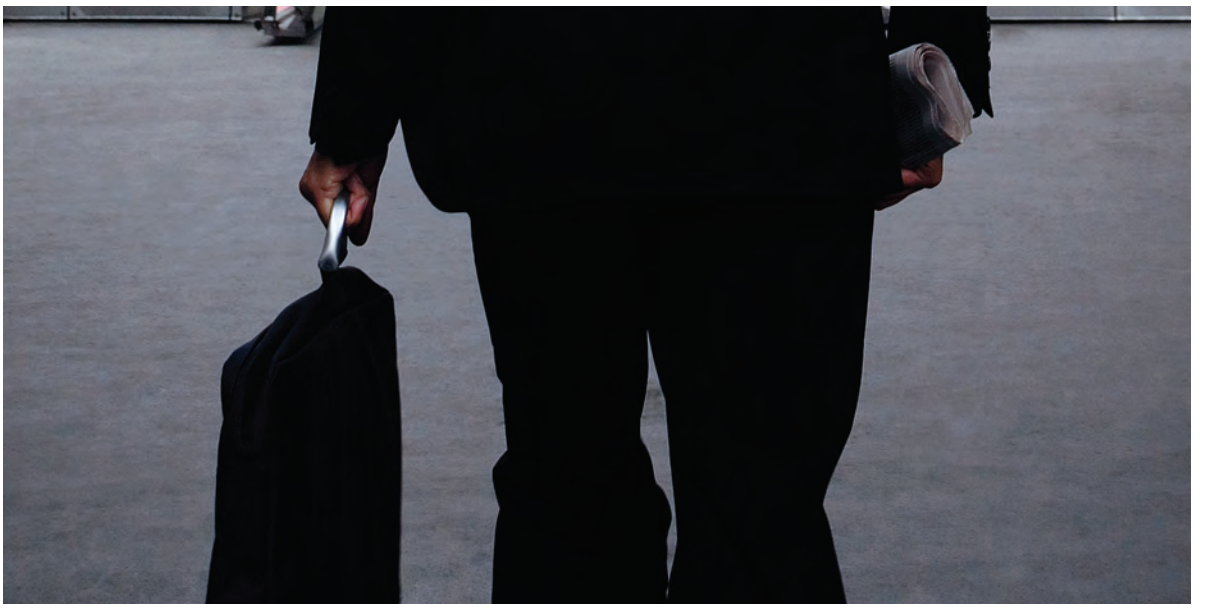

Security Handbook Contents

HP recognizes the complexity of large, distributed IT environments and takes a proactive approach to enterprise security. We secure the Adaptive Enterprise with planning and preparation, rather than simply reacting to changes in the landscape. This handbook outlines HP's strategy for information security and summarizes the products, solutions, and services that address the security needs of enterprise customers. It focuses on the three pillars of HP's security strategy: proactive security management, identity management, and trusted infrastructures. Along with overarching governance considerations, these three areas bring organizations and companies a safer IT environment that can respond to changing threats and business objectives.

This edition of the handbook introduces a section on innovation in information security pursued at HP Laboratories, HP's central research organization. Uncoupled from product and services organizations, HP Labs' mission is to deliver breakthrough technologies that create opportunity beyond HP's current strategies. Some of the work performed at HP Labs has generated new capabilities which are reported throughout the handbook. This chapter addresses aspects of the longer-term challenge in information security and the work pursued by HP Labs to overcome them.

This handbook is intended for CIOs, security administrators, and other staff who are responsible for their organization's IT security and infrastructure. Each chapter begins with the definition and purpose of the topic before moving on to discuss details such as the threat environment, related trends, underlying technologies, and challenges. Each chapter concludes with information about solutions that address the security needs discussed.

Bringing the HP Security Strategy to Market: HP Secure Advantage

Security today is higher on the CIO/CTO agenda than it ever was before. They believe that security is a fundamental necessity and enabler of business outcomes. They want to be able to use HP products and plug them together easily in a secure manner. HP Secure Advantage is the framework under which this will take place. Imagine data protected from desktop to data center, from laptop to printer, throughout the network with no gaps; no places where the data has to be decrypted and re-encrypted to transition to another product. Imagine being able to demonstrate to internal and external auditors that you have a trusted infrastructure and that your data in any form is protected so that you can easily add the people processes to meet your compliance demands; that is the promise of Secure Advantage.

As this Security Handbook illustrates, HP has a unique breadth of products, from laptops to servers and storage and software to printers, using HP network components. This is why HP created the Secure Advantage framework and a portfolio of products and services to meet our customer's needs for secure data and infrastructure protection. Fortunately, HP has a 35 year history in security and is leveraging this expertise to deliver the HP Secure Advantage portfolio. This is especially important today as customers adopt the 24 by 7 next-generation data center model that enables the shift of high-cost IT silos to low-cost, pooled IT assets in order to optimize infrastructures to reduce cost, increase agility, and improve quality of service. Security is a key enabler of HP's Adaptive Infrastructure (AI) offering that provides the platform for the next-generation data center and a linkage of Security to other AI enablers such as IT Systems and Services, Power and Cooling, Management, Virtualization and Automation.

Since security and compliance are an absolute necessity for businesses today the HP Secure Advantage portfolio is designed to enable enterprises to fully automate, optimize and accelerate their IT infrastructures securely with proper validation in order to achieve better business outcomes by mitigating risk. In order to accomplish this goal across the Enterprise, HP is establishing leadership for solutions in information security, key management and compliance.

The HP Secure Advantage vision builds on today's security technologies to create a more manageable way for customers to leverage encryption and key management to protect their resources and data, and validate they are compliant with a growing set of government and industry mandates.

The HP Secure Advantage portfolio takes a layered and integrated approach and helps you extend the value of your enterprise by the following steps - as illustrated in Figure i-3 and explained in the following paragraphs:

Figure i- 3
HP Secure Advantage overview

Protect your resources by improving availability and protecting your networks, systems, applications, software and DBMS using trusted platforms. Minimize IT disruptions due to security breaches with a trusted and hardened infrastructure:

- Multiple OS platforms with the highest level of certification provide maximum proactive protection.
- Configuration and patch management provide continuous protection in changing environment.

Protect your data in all its forms: data at rest, data in transit, and data in use. Use encryption and Identity Management in combination with other proactive security management techniques such as Security Event and Information Management:

- Encryption of critical data at rest, in use or in motion increases protection.
- HP Extends this protection from desktops to servers and printers with focused Key Management.

Provide validation to you by establishing a secure audit trail across the organization as proof for compliance for internal and external auditors with real-time alerts. Utilize encryption and Key Management, working with integrated compliance solutions across organizations.

- Validate at necessary audit points to enable audit trails for compliance to industry regulations.
- Future integration of encryption and Key Management across an organization will provide end-to-end protection.

These elements are more integrated under the Secure Advantage portfolio framework and can be customized to your unique needs by HP Services through our Information Security Service Management (ISSM) Reference Model which establishes a rational basis for security decision-making, ensuring security controls align, and helping optimize business outcomes.

Table i-1 provides a set of examples of key HP Secure Advantage products and services. The table also shows the corresponding HP Security Handbook chapters, sections and pages that provide more detail on these HP Security Advantage offerings.

Table i-1
HP Secure Advantage products and services examples and HP Security Handbook links

HP Secure Advantage Product/Service	HP Security Handbook Link
Protect your Resources	
HP Configuration Management	Chapter 2 - Proactive Security Management: Section 4.2.3.3.1, page 2-23
HP Proliant Essentials Vulnerability and Patch Management Pack	Chapter 2 - Proactive Security Management: Section 4.2.3.3.2, page 2-24
Protect your Data	
HP Compliance Log Warehouse	Chapter 4 - Trusted Infrastructure: Section 7.4.3, page 4-73
HP Secure Print Advantage (SPA)	Chapter 4 - Trusted Infrastructure: Section 8.5, page 4-78
HP StorageWorks Secure Key Manager	Chapter 4 - Trusted Infrastructure: Section 7.4.2, page 4-72
HP ProtectTools	Chapter 3 - Identity Management: Section 8.2.2.1, page 3-16 Chapter 4 - Trusted Infrastructure: Section 5.4.1, page 4-28
HP- UX 11i security	Chapter 4 - Trusted Infrastructure: Section 5.4.3.1, page 4-31
Linux security enhancements	Chapter 4 - Trusted Infrastructure: Section 5.4.3.3, page 4-36
HP NetTop	Chapter 4 - Trusted Infrastructure: Section 5.4.2, page 4-30
HP Application Security Center	Chapter 2 - Proactive Security Management: Section 4.2, page 2-19 Chapter 4 - Trusted Infrastructure: Section 5.4.6, page 4-42
HP ProCurve Identity Driven Manager	Chapter 4 - Trusted Infrastructure: Section 6.8.2.2, page 4-65
HP Trusted Compliance Solution for Energy (TCS/e)	Chapter 4 - Trusted Infrastructure: Section 5.4.5, page 4-42
Provide Validation	
HP Services Information Security Service Management (ISSM)	Chapter 1 - Governance and Compliance: Section 8, page 1-16

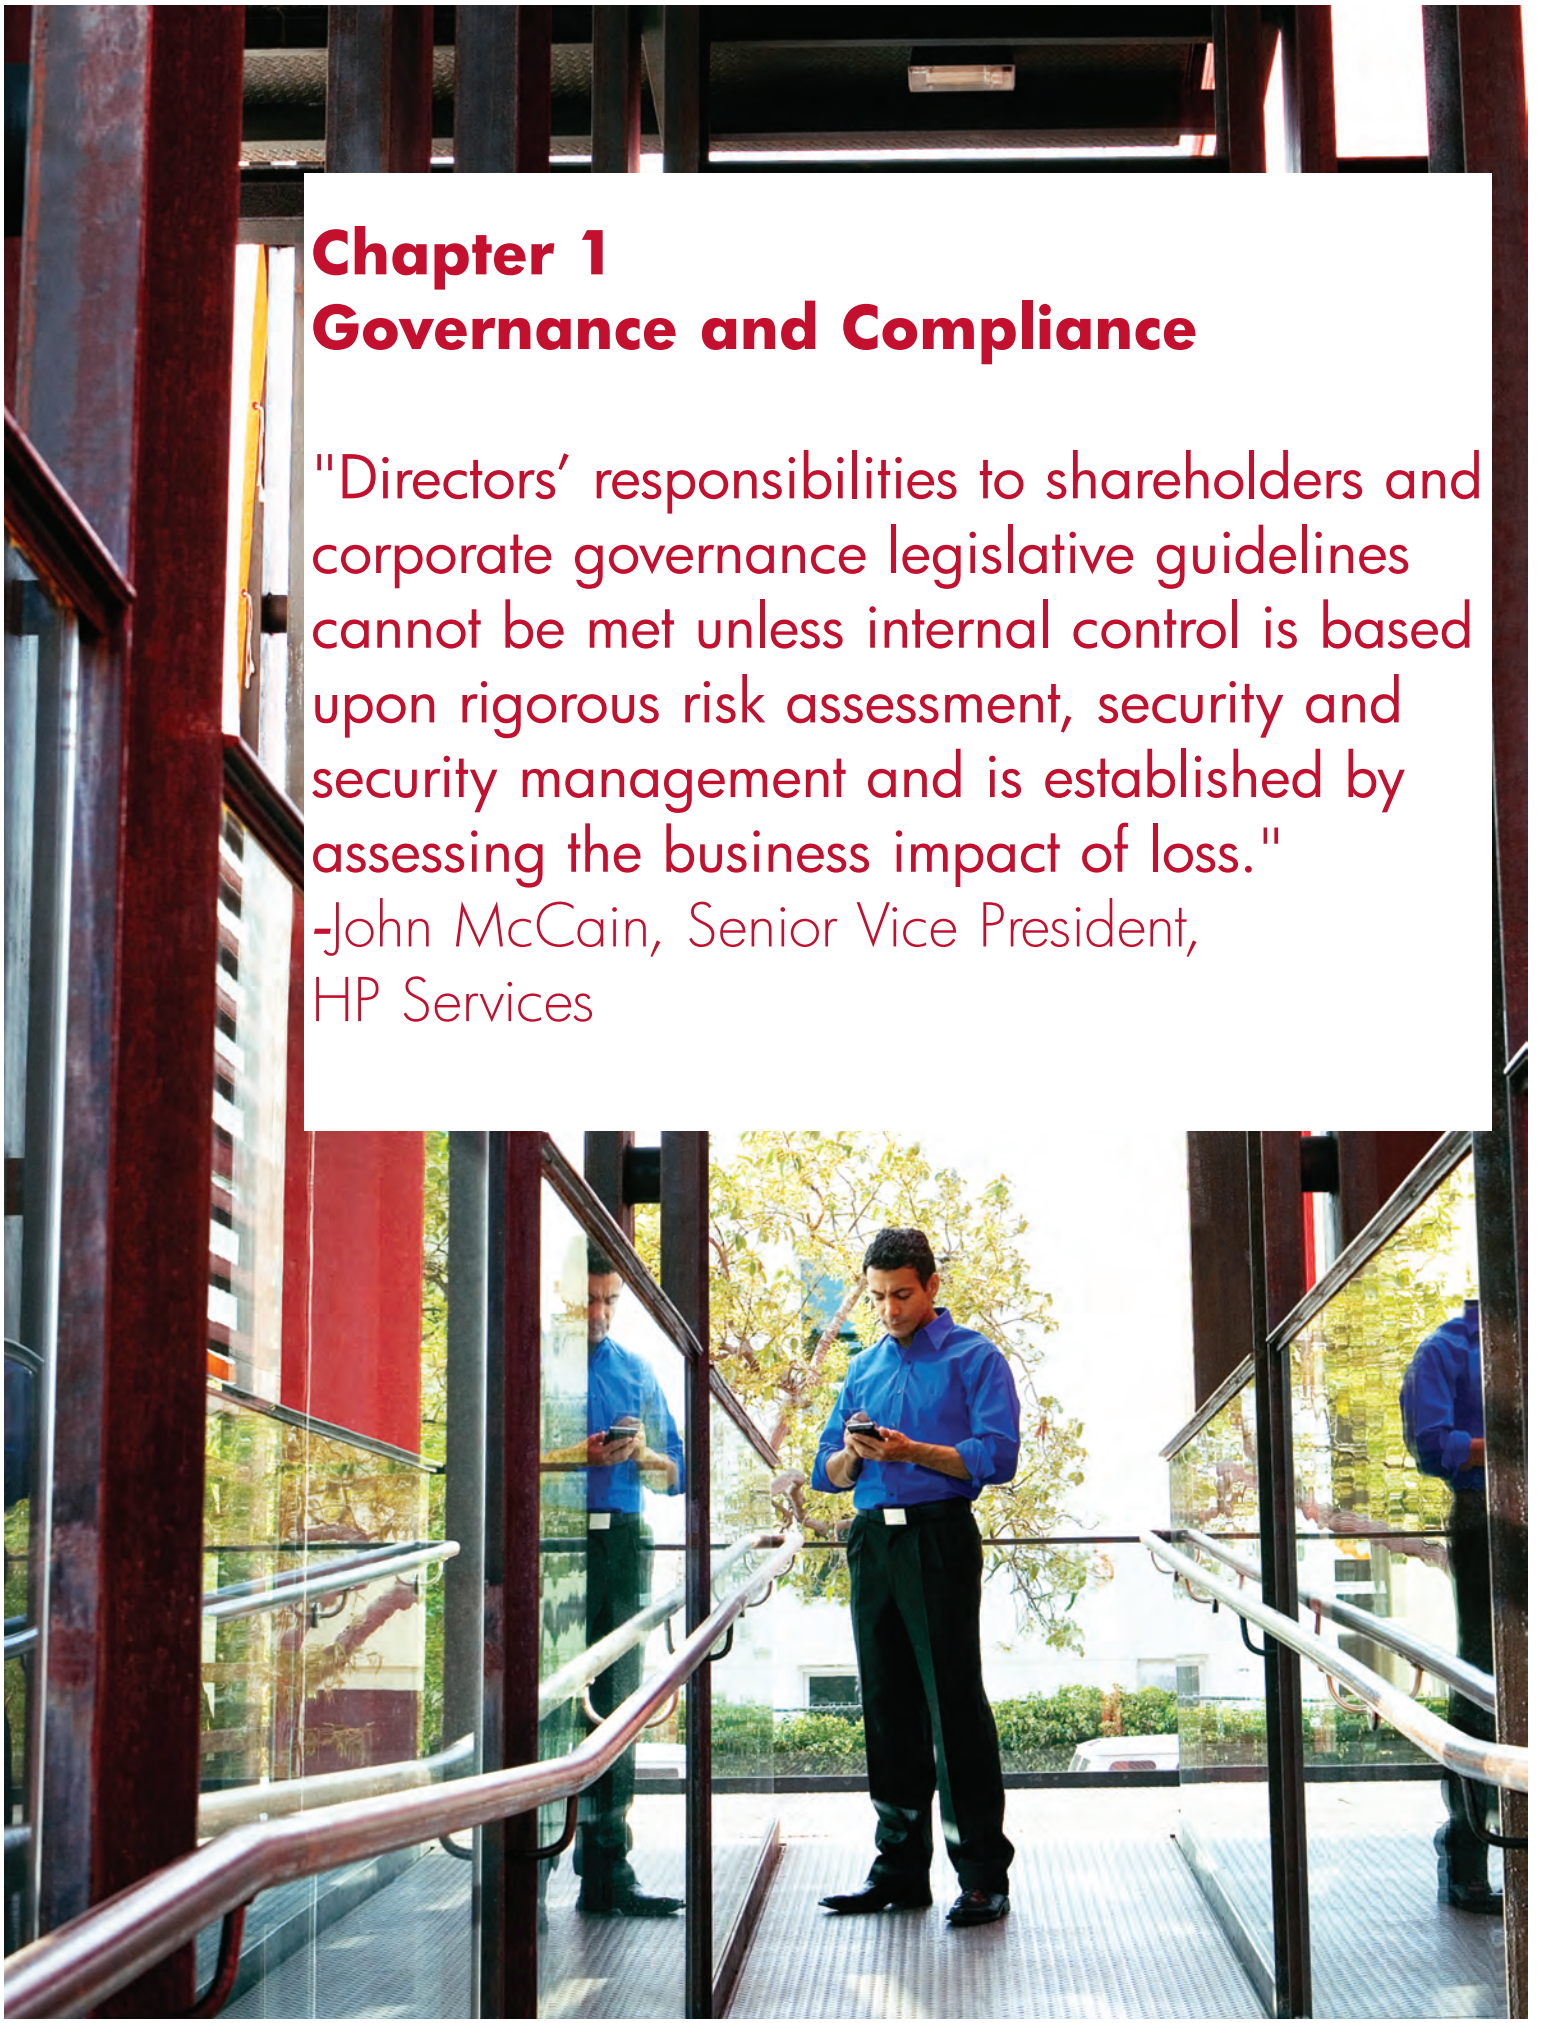

Chapter 1
Governance and Compliance

"Directors' responsibilities to shareholders and corporate governance legislative guidelines cannot be met unless internal control is based upon rigorous risk assessment, security and security management and is established by assessing the business impact of loss."
-John McCain, Senior Vice President,
HP Services

Figure 1-1
Governance and Compliance

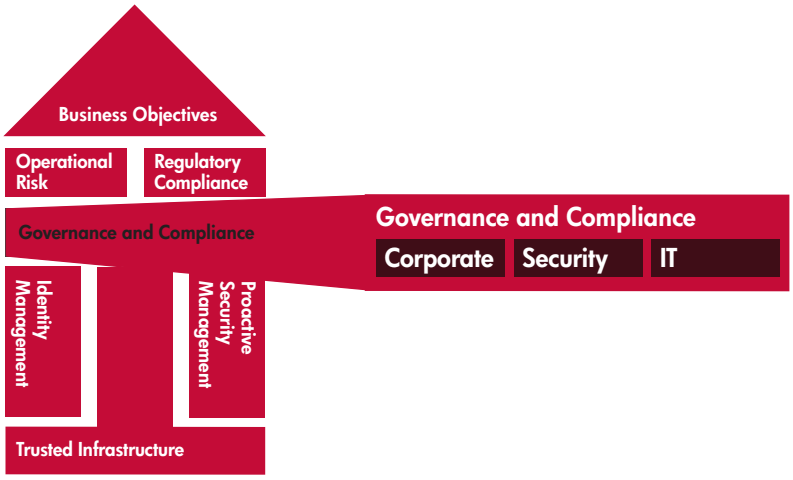

"Due diligence is the objective, governance is the process used to achieve the objective and compliance is the way you measure achievement of the objective."
-Stuart Hotchkiss, HPS Consulting and Integration

The role of governance has never been more critical. Corporate leaders are required to move from a best-practices approach to a legislated approach to asset protection, and IT departments are firstly charged with delivering business value while also reducing operational risk through the establishment of effective business controls, all while providing continuous compliance, state attestation and reporting. HP recognizes this dilemma and has developed a security framework for governance that provides the overall structure for implementing business objectives, complying with regulations, adhering to risk strategies, and protecting information assets. Security governance supplies a critical link between business management and IT.

This chapter begins by defining governance and the responsibilities of company officers in meeting legislative requirements. It details the recommended governance lifecycle and the steps that need to occur to achieve the objectives of compliance with business and regulatory requirements. The chapter ends with a review of how to move from a static to an operational view of governance and continuous compliance.

1. Definition

Governance refers to the controls and policies that translate high-level business objectives, operational risks, and regulatory needs into the directives, objectives, and policies that drive security mechanisms.

Governance is a strategic component of every business technology optimization initiative. It contains business logic, business procedures, and managerial and operational processes, which are supported by more specific, lower-level policies for IT operations and security. Governance is often classified and tiered as corporate governance, IT governance, and security governance.

Corporate governance is the process by which a company's board of directors achieves two objectives for shareholders:

- The efficient use of business assets and resources
- The availability of assets for new business to maximize shareholder return

Corporate governance achieves these objectives by defining the risk profile of the company's businesses and investing in appropriate controls to allow the organization to function effectively with minimal operational disruptions within their regulated requirements. It explicitly defines the risk appetite of the enterprise and the mitigation methods for anticipated and unforeseen risks. These mitigation methods are developed by functional staff (including IT), but the board of directors should review, understand, and agree to them.

Various regulations formalize corporate governance, either specifying adherence to the recognized industry best practices or dictating a very prescriptive process, such as Sarbanes-Oxley. However, complying with regulations does not guarantee good governance. Noncompliance with regulations, on the other hand, is a fairly sure sign of poor governance.

Figure 1-2
Governance overlaps

The intent of most regulatory acts such as Sarbanes-Oxley is to remove subjectivity from governance. In the past, codes of practice and best practices have been common methods for demonstrating suitable governance and auditing has been used to prove control. This has been acceptable as long as the results were satisfactory and fraud was not proven. However, regulatory compliance is often perceived as the most important issue. This masks the fact that regulatory compliance is designed to be an indicator of good governance - not good governance in and of itself.

Security governance is a component of corporate governance. It is the requirement of company directors to demonstrate due diligence in handling information assets on behalf of stakeholders. Security governance is composed of all the processes and decisions that affect company assets in terms of their validity, confidentiality, integrity, and availability for business. Without security governance, corporate governance objectives cannot be met simply because there can be little faith in the internal control systems.

In this context, security governance encompasses all assets and their threats. Therefore, physical building security and transport security, for example, are part of this process. When assessing risk, the threats to all assets should be reviewed. Some mitigation plans, consequently, will include IT components and some will not. In this sense, security governance is wider in scope than IT governance, as shown in Figure 1-2.

As with corporate governance, mitigation plans for security governance fall within the bounds of the risk appetite as expressed by the board of directors in the company objectives. Likewise, the board should explicitly approve the final choice of mitigation plans and the controls for each.

IT governance ensures that IT supports business requirements and that it does so efficiently and flexibly. This subject exists simply because IT and business often misunderstand each other. In particular, the differing time scales, language, priorities, expectations, and contexts of IT and business can lead to a disconnect. The IT Governance Institute (**www.itgi.org**) is a good source of information about best practices for IT governance and its alignment with security governance. IT governance is directly affected by security governance; IT cannot produce reliable results if security is inadequate.

IT systems, technologies, and processes are at the core of most businesses. As such, they have two important roles: to facilitate business efficiency and to mitigate business risks by implementing controls. Because they have differing impacts, it is important to understand these IT roles.

2. Purpose

Shareholders and legislators can require directors to prove that they have taken due care in the use of company assets. A company asset is anything that the company owns and uses in business - from office chairs to information. A lack of security (IT or physical) is usually caused by a lack of due care. To make it easier for company directors to prove due care, every country has a number of legislative frameworks for directors to adhere to - some are statements of self-regulated best practices and others are very prescriptive. If due care cannot be demonstrated, directors can be removed from office, fined, imprisoned, and subjected to the loss of personal property. This occurs in every country in the world today, and it came about long before the Enron scandal.

3. Why Have Governance?

Governance regulations and guidelines attempt to provide a prescriptive management framework and an independent method of determining how well businesses are managed. In addition, there is an increasing need for companies to be comparable worldwide. The various stakeholders in a company have differing needs related to governance:

• Shareholders and regulators look for adherence to corporate governance control frameworks and regulations to determine how well the company is managed.
• Auditors and regulators refer to demonstrated adherence to security guidelines and control systems to determine whether the company applies basic due diligence and control.
• Management follows a security management lifecycle to ensure that business controls and IT governance needs are met.

4. The Importance of Information Assets

One of the most sensitive areas of modern business is the exposure faced by information assets. In some cases, these assets make up the majority of a company's capital, and their loss or damage can put a company out of business. A graphic example is a credit card company that exposes all of its credit card details such that massive fraud can occur. Such an event exposes the company to legal liability, lost clients, and damages. It would undoubtedly shut the company down and lead to the sanctioning of the company directors for dereliction of duty.

Security governance is comprised of all the actions that directors need to undertake to avoid such events and to prove to authorities, business partners, staff, shareholders, and clients that they are treating company assets in a secure manner. From another viewpoint, if information assets are insecure, IT cannot produce reliable results. Therefore, company directors cannot report accurately or manage assets correctly. Reporting accurately is a key component of Sarbanes-Oxley, the Basel II Accord, and most corporate control frameworks.

In this context, the scope of security governance is the security of operations relating to end products. It is not the security inherent to the product in and of itself. In a financial services business, for example, a credit instrument has the risk that the client will default. This risk falls under corporate governance. However, losing the credit instrument's details and falsifying transactions are risks within the scope of security governance.

5. Information Security Defined

Information security refers to the security of information assets. The most widely used characteristics of information security are confidentiality, integrity, and availability:

• Confidentiality means that only the user or the user's delegates have access to the information.
• Integrity means that the information is in the expected state and that it has not been changed without knowledge or permission.
• Availability means having the information when it is needed.

In reality, availability has the greatest impact. If information is not available and business must continue, is outdated information used instead? Other characteristics that are commonly referenced include utility (whether information is in a useful state) and non-repudiation (someone who uses the information cannot deny it later). These characteristics are best viewed as consequences or applications of the three basic characteristics. For example, non-repudiation requires an application to use information that can remain confidential, unchanged, and available in a useful state.

When securing the confidentiality, integrity, and availability of information assets, organizations should examine the entire information environment. This key part of information security is often forgotten. It is common for the IT part of the security equation to be separated financially and organizationally from the wider information security environment, often with dire consequences. For example, information security cannot be guaranteed if:

• Information handling processes are not defined, including backup, restore, and off-site storage procedures for sensitive or critical data.
• Audits are not possible because a reliable current state does not exist.
• Information-processing facilities do not have adequate physical security and protection, such as appropriate fire suppression systems.
• Persons transporting (or with access to) systems or data have not been vetted.
• Information lifecycle policies and procedures -from creation to destruction - do not exist.
• People, processes and technologies do not work in concert.

These are only examples, and a framework for information security covers most areas. The key is to examine threats to the entire environment and assign mitigation methods and sufficient resources to the totality of the risks. This occurs during the risk assessment process, which should be ongoing.

5.1. Board of Directors' Responsibilities

Within all aspects of governance, a company's board of directors is responsible for explicitly defining aspects of control and security. Legislation such as Sarbanes-Oxley requires explicit sign-off. This requirement, however, has existed in the compliance frameworks of most countries for many years, but it was rarely enforced.

Specifically, the governance responsibilities of the board of directors include:

- Understanding the subject of information security in general
- Setting direction and driving policy and strategy
- Defining and agreeing to risk appetite for defined businesses and reviewing and accepting risk mitigation proposals
- Providing resources
- Ensuring that individual roles, responsibilities, and authority are clearly communicated
- Delegating authority, but not responsibility
- Approving security measures explicitly
- Reviewing risk appetite and security measures periodically
- Implementing an organization that enables security governance, with the security department reporting directly to the board and not another business function

The board can delegate these responsibilities, but it is held accountable. For example, outsourcing IT or asset management does not alter the board's ultimate responsibility for these functions.

5.2. IT Responsibilities

By the same token, IT has a number of responsibilities. Its key function is to implement business controls. For this reason, IT should aim to fully understand business objectives. If management and IT are not in alignment (and IT does not fully understand and support the relevant business drivers and priorities), IT will not have the appropriate context in which to frame and plan security architectures.
Specific IT responsibilities for governance include:

- Participating in business impact analysis exercises with business managers
- Proposing and gaining agreement for risk mitigation strategies
- Developing architectures that implement current and potential control requirements
- Identifying threats and analyzing vulnerabilities in proposed technical components, and staying current with updates, patches, and new threat vectors
- Implementing monitoring and incident response methods
- Conducting periodic reviews with audits
- Ensuring application security during development and when acquiring applications externally
- Ensuring awareness of the need to protect information and recommending relevant user training.

Due to differing time scales or rapidly changing needs, IT projects may not synchronize with business requirements. Within this context, IT has the responsibility to inform the board of the limits and capabilities of IT and to help improve efficiency in line with business objectives.

6. Regulatory Standards

Every country has legislation concerning the requirements of directors with regard to due diligence and governance. Although some legislation and standards are very specific, the majority are non-prescriptive guidelines. As a result, most regulatory standards outline directors' responsibilities without explaining how to accomplish them.

A common theme is an external audit of a control system. The audit is usually of a financial nature to ensure the accuracy of the numbers reported. However, using audit alone as a tool to improve a control system is probably not a good idea. Accounting standards and regulations can be applied worldwide, for example, but they are influenced by local government tax policies. This makes it difficult to compare results.

Another issue with using audit only is the cost associated with control system failures. The cost can be accounted for in many ways, and the manner of accounting can impact an audit. For example, in the past, banks commonly accounted for these failures as operations costs and passed them on. Now banks have a clear obligation to separately report any control system failures that lead to operational losses, in a way that penalizes the bank.

A final difficulty is that audit methods usually audit against a known or required state. If this state is not defined correctly in the first place, audit is a weak tool for governance. Sarbanes-Oxley is the first framework and legislation that provides enough information for companies to establish a clear control system and track its performance. This is evidenced by older control frameworks such as SAS No. 70 where organizations audit themselves against their own declared state.

Faced with these difficulties, the best course of action is to improve control systems from the bottom up. Conforming to international standards addresses the majority of control system problems. Organizations can then focus on the few remaining concerns.

6.1. International Standards

Examples that organizations use internationally to demonstrate or regulate governance include: the International Organization for Standardization (ISO), the Control Objectives for Information and Related Technology (COBIT), and the Committee of Sponsoring Organizations of the Treadway Commission (COSO).

6.1.1. The ISO/IEC 27000 Standard Family

This is the generic name for a whole series of intended information security related standards. The ISO 27001 standard was published in October 2005, essentially replacing the old BS 7799-2 standard. It is the specification for an ISMS, an Information Security Management System. BS 7799 itself was a long standing standard, first published in the nineties as a code of practice.

The ISO 27002 standard is the renaming of the existing ISO 17799 standard, ISO 17799 being a code of practice for information security.

ISO 27003 is a proposed development to provide help and guidance in implementing an ISMS. This will include focus upon the PDCA method, with respect to establishing, implementing reviewing and improving the ISMS itself.

As an example, best practices state that a firewall should protect a network perimeter and that a business continuity plan (BCP) should address information availability. However, a code of practice does not address how to manage the firewall, when to update the firewall, or what process to use. Similarly, it does not define how to develop or manage a BCP. The reason for this is clear - each company has its own view about the assets to protect and the acceptable risks. ISO 27001 gives a specification of a management system which can be used. Although, like all standards, it can never define exactly what to do for the risk appetite of a particular company since this is a management judgment.

6.1.2. Control Objectives for Information and Related Technology (COBIT)

COBIT was developed by the Information Systems Audit and Control Foundation (www.isaca.org) and the IT Governance Institute (www.itgi.org). It provides a broad control framework together with control objectives for all of IT. COBIT splits IT into 34 process areas within four domains. There is a four-step cycle from planning through monitoring, and the goal is to provide a control framework that allows management to audit and evaluate how well IT processes align with business.

Only a portion of COBIT focuses on information security, however. It also provides an excellent means of determining the degree of maturity of the control systems within a company. Like ISO 17799, it defines what to do, but not how to do it. COBIT adds the concepts of authenticity and non-repudiation as base requirements within information security, and it uses seven information criteria to define what business requires from IT:

- Effectiveness
- Efficiency
- Availability
- Integrity
- Confidentiality
- Reliability
- Compliance

6.1.3. Committee of Sponsoring Organizations of the Treadway Commission (COSO)

The COSO report defines internal control as a process, affected by an entity's board of directors, management, and other personnel. It is designed to provide reasonable assurance regarding the achievement of objectives in the following categories:

- Effectiveness and efficiency of operations
- Reliability of financial reporting
- Compliance with applicable laws and regulations

The COSO framework is mentioned in the Sarbanes-Oxley proposals as a framework suitable to satisfy its requirements, although others such as COBIT also qualify.

The original version of COSO (Internal Control-Integrated Framework) has been augmented by the addition of the Enterprise Risk Management-Integrated Framework. It uses the same conceptual foundations but offers a broad, risk-based approach to strategic control. For more details, see **www.theiia.org**.

6.2. When to Use the International Standards

6.2.1. ISO 2700x Series

ISO 27002 defines best practices in fairly clear terms and covers ten main areas of business. Following this standard ensures that all areas of a business are reviewed. However, it does not ensure that security is good or sufficient. An assessment against this best-practice framework covers the following ten areas:

- Security policy
- Organization of information security
- Asset management
- Human resource security
- Physical and environmental security
- Communications and operations management
- Access control
- Information systems acquisition, development, and maintenance
- Business continuity management
- Compliance

ISO 27001 defines a management system for security. An external audit can use this framework, but there are important comparability issues. For example, the scope defines which part of a company is reviewed. Unless two companies use an identical scope, there is no comparability. In addition, the Statement of Applicability (SOA) is the management assessment of whether a control is necessary or sufficient. Unless the SOAs are identical, comparability suffers.

6.2.2. COBIT

This framework is ideal to review the needs and responsibilities in IT governance from a wider perspective than a pure security view. It takes the position that there are information requirements that align IT with business. The starting point is defining an IT strategic plan. Next, the plan is detailed and the necessary components are acquired and implemented. Finally, the plan is delivered, supported, and monitored.

Handling risk is an implicit understanding within the COBIT framework. However, it does not specifically examine the risk to an information asset and the subsequent mitigation process used by IT. For example, COBIT does not examine the risk of losing an invoicing system at a cost of $10,000 per day or the resulting manual and IT mitigation process. Instead, COBIT examines the risk inherent in the IT methods, and using COBIT, a business owner can be sure that an adequate, efficient control system is in place for the IT environment. Note that COBIT is weaker in general security and the process side of IT.

It is often helpful to use the control objectives of COBIT, but to report them within the framework structure of ISO 2700x. For more information about COBIT, visit **www.isaca.org**.

6.2.3. COSO

COSO is a comprehensive and fundamental internal control framework. As such, it is an excellent base for corporate governance. For real applicability, however, more specific frameworks and standards need to be used, such as COBIT, ISO 2700x or other best practices that consider the impact of a control. As a general point, it is important to define the impact of a control and determine whether to use it based on the risk it helps mitigate.

Sarbanes-Oxley mentions COSO, but it is not a prerequisite for compliance. Any control system that demonstrates reliable results and follows best practices is suitable. The COSO control framework includes:

- Control environment
- Risk assessment
- Control activities
    - General controls (data center and software access)
    - Application controls (development methodologies and application controls like checksums)
- Information and communication
    - All types, but some focus on control information reports
- Monitoring
    - Continuous
    - Point

COSO is not sufficiently prescriptive to handle security in the broadest sense. There is a strong emphasis on the organizational and cultural requirements to embed risk management and control into a company, but the strong and directive links between risk, risk mitigation, and the mechanisms required in both IT and general security controls have less emphasis. Failure of any part of the underlying security mechanisms that ensure data quality can invalidate the entire control system. A concerted effort is needed to ensure the confidentiality, integrity, and availability of all information assets involved in this equation. For more information about COSO, visit **www.coso.org**.

6.3. Best-Practice Legislation

The most effective way to ensure that stakeholders such as legislators, auditors, clients, business partners, and staff recognize security governance is to comply with best-practice legislation. Many legislative texts affect information security, but best-practice legislation requires a complete security governance program for reliable and permanent reporting to regulators.

There are a number of important legislative frameworks dealing with different domains and applications, with common threads in terms of controls and best practices:

- Sarbanes-Oxley Act of 2002
- Basel II Accord
- Health Insurance Portability and Accountability Act (HIPAA)
- Gramm-Leach-Bliley Act (GLB)

These texts describe best practices and provide guidelines; however, they do not specify how to comply with them. As an example, the Basel II Accord clearly states the need for business continuity, but it does not provide details or propose the use of a standard body of knowledge, such as the Business Continuity Institute or DRI International.
Faced with this, it is difficult to know how to satisfy regulators. However, as a basis for compliance with best-practice legislation, aiming to comply with the provisions of ISO 2700x and COBIT provides the foundation for compliance with all regulations and for alignment of IT security with business needs.

6.4. Privacy Aspects and Issues

Privacy management is a core concern for enterprises and people, and it requires integration with governance efforts. From an enterprise perspective, privacy management is a necessary aspect of regulatory compliance because governments and corporate guidelines require it. Regulatory laws such as Sarbanes-Oxley, GLB, HIPAA, and various governmental directives on data protection require enterprises to implement complex processes to comply with related policies.

Specifically, much work has been done in terms of privacy legislation, often driven by local or geographical needs. This includes European Community data-protection privacy laws, various U.S. privacy laws, and more specific national privacy initiatives. Guidelines are also available for protecting privacy and the flow of personal data, including the Organization for Economic Co-operation and Development (OECD) guidelines. The OECD guidelines describe concepts such as collection limitation, data quality, purpose specification principles, and online privacy policies. For more information about the OECD, visit **www.oecd.org**.

Enterprises store large amounts of personal (confidential) data about their employees, customers, and partners. Failure to comply with privacy policies can result in serious consequences for the reputation and brand of organizations as well as negative legal and financial impacts. Furthermore, a large enterprise with a multi-national presence might need to comply with international privacy laws. Additional nations (such as South America and Japan, where privacy law came into effect in 2005) are developing privacy legislation.

Privacy legislation contains common provisions for securely handling, storing, and disseminating information. The same underlying requirements of due diligence for security governance provide the right framework of proof for this aspect of privacy legislation. Because privacy legislation is more closely monitored than governance regulations (any client, employee, or supplier can file a complaint), these security governance practices become mandatory.

7. The Governance and Risk Management Lifecycles

Ideally, business needs drive the governance lifecycle. There are many key decisions that only the board of directors or its delegates can make. The board needs to define objectives, identify the risk appetite, and explicitly sign off on mitigation plans for business risks. It must also be aware of the operations of security implementations in terms of how the operations affect governance needs.

It is common for the process of security to be technically oriented and not driven from a business perspective. This causes two main problems. First, the security equation focuses on technical assets, rather than all the assets required to conduct business. The second problem is that the analysis often examines the cost of mitigation methods (and usually chooses the least expensive) without balancing the cost with the potential business impact. This lack of synchronization often results from the perception that security is an IT problem or from a communication/organizational gap between business and IT. The need to provide evidence of governance and due diligence via this process is a major requirement.

Another common oversight is to drive the process of security with security policy. This is a partial mistake. Security policy is derived from an analysis of business needs, and policies are only one of the tools necessary to mitigate risks. Policy should be firmly placed within the mitigation area - if a business asset is of no value, then policies do not need to protect it. Business analysis must be performed to determine how much to spend on risk mitigation, because it corresponds to the business value of the resources and assets being protected.

Risk mitigation is driven by and accountable to the governance bodies of an organization. The governance bodies function in a two-tier system. First, the corporate and business governance bodies define the level of risk they are willing to accept. Second, the IT and risk management governance bodies take this input and interpret the objectives. Figure 1-3 illustrates the relationship of governance, risk management, and daily security operations.

Figure 1-3
The governance lifecycle

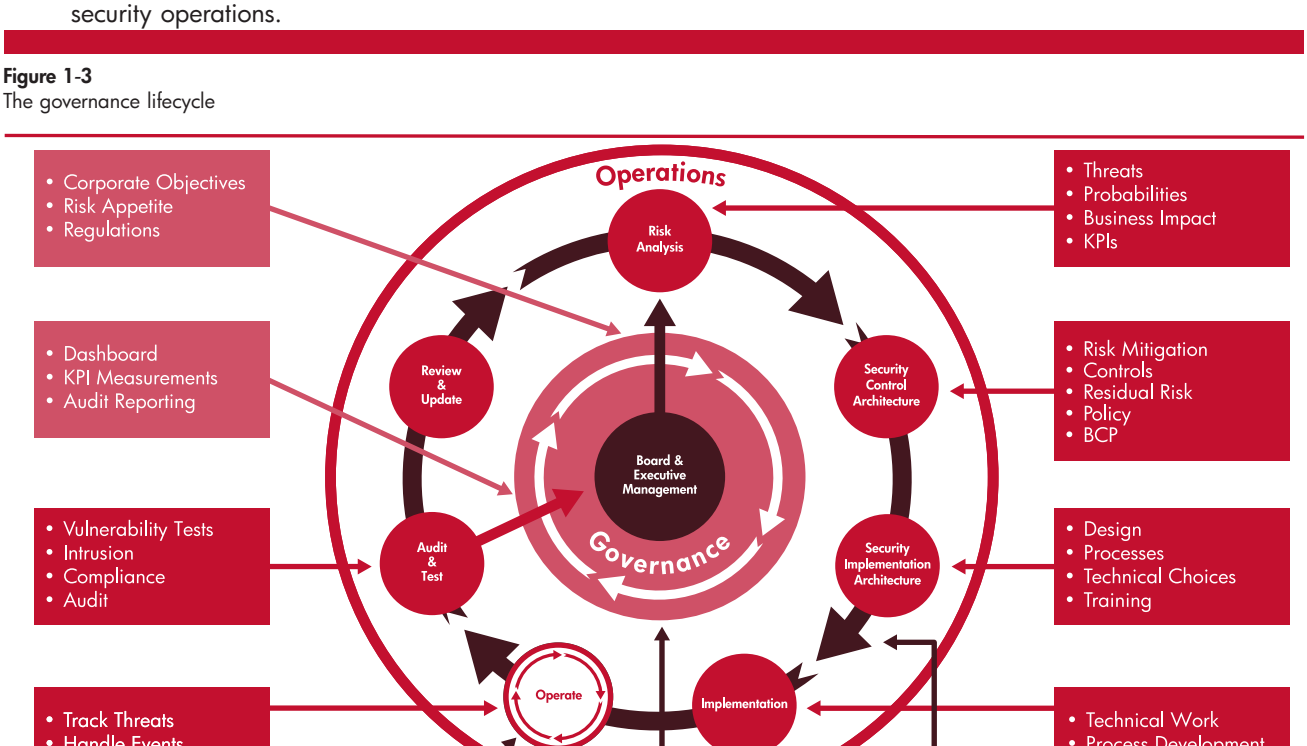

The information security governance lifecycle corresponds to the processes that operate across the governance areas of corporate, business, IT, risk management, and information security. The purpose of this lifecycle is to allow executive officers to exercise their risk management oversight responsibilities. It interfaces with the risk mitigation lifecycle by setting the objectives for risk mitigation, including the definition of risk appetite. Additionally, it interfaces with the risk mitigation lifecycle by consuming dashboards, Key Performance Indicators (KPIs), and audit reports. Beyond these specific interfaces, the two groups interact regularly on a more or less formal basis.

The risk mitigation lifecycle transforms the risk management objectives set by the governance bodies into controls. It implements these controls and regularly reviews their effectiveness. This lifecycle clearly separates identification of risks (based on business impact, not technology), identification of controls, design of control implementation, and implementation of controls through technology and process. This lifecycle is typically iterated once a year.

Daily security operations function continuously; this is where the IT system delivers its services to the business. This set of processes produces a massive volume of event data. The data is the base material for security audits and reports to the governing bodies.

These lifecycles take place on the background of operations, where mitigation methods are executed. The majority of these methods use IT technologies and processes. IT is responsible for efficiently implementing and managing operations, in line with IT governance principles. IT does not have sole responsibility for any of the other process steps. However, IT is responsible for providing input and feedback with respect to relevant policies and operations (where appropriate). Continuous open communication and dialogue are keys to ensuring IT alignment with process goals.

7.1. Process Steps

In a new company, the ideal first process step is to perform a business risk analysis. Subsequent steps might rely on existing architectures and may cause them to be modified. If a business has new risks, for example, it may require new policies or technologies that modify the existing security architecture. The term security architecture as used here implies all of the technology, people, and processes required to implement the security governance objectives of a business and to align these objectives with the requirements of IT governance, due diligence, and legislative compliance.

Existing organizations may choose to start elsewhere in the cycle or with a gap analysis. However, at some stage they must align with the cycle so that the driver is business impact analysis (BIA). It is simply not possible to determine how much money to spend or which mitigation methods are suitable until the potential loss is assessed. The most common mistake, which leads to overspending or underspending, is to drive the lifecycle via technology.

Note that the lifecycle presented here is a general guide; actions can vary and the content of the steps should be tailored to each organization. Some principles do apply; for example, good governance requires monitoring, a minimum yearly review of risks, and control systems to mitigate those risks.ife

7.2. Gap Analysis

A business that does not have a security architecture or controls is usually characterized by regular spending on a yearly or project cycle basis. From the outside, this looks like a set of remedial actions, rather than a coherent plan. In this case, it is best to start the lifecycle by performing a gap analysis, which highlights the differences between the current state and the desired future state.

Many events can trigger the need to do a gap analysis. For example:

- An audit reveals that policy was not followed.
- An intrusion reveals that the network was badly configured.
- A business outage reveals that continuity plan training was not followed properly.

It is not necessary for an event to trigger a gap analysis. A desire to move to a standard can begin a gap analysis against that standard. In addition, a gap analysis can occur at any point during the lifecycle - it is not necessarily at the beginning.

Documentation of the current expected state is the critical issue for a gap analysis. Gap analyses often show that the expected state was not documented, or it was well documented some years ago and has evolved since. A simple example is to define and document the current expected working environment of the average PC.

There are many types of gap analyses, including:

- *Standards framework*: In this case, the organization's desired state is usually a framework such as ISO 27001, COBIT, or COSO. A standards framework measures the state of conformance to the standard. Defining or documenting a current state in order to measure the conformance is not necessary.
- *Policy framework*: Within this framework, an organization's policies, standards, and guidelines are defined and documented. The policy framework measures their implementation.
- *Audit results*: After an audit, the lifecycle process begins by correcting problems or issues identified in the audit results. The audit can be of any kind, including an external audit, a vulnerability test audit, or an intrusion test audit.

7.3. Risk Analysis

Risk analysis is the key process step for the governance lifecycle. As the normal starting point, risk analysis identifies all business processes and documents the applications and people processes for each. Business management performs the risk analysis to ensure that the perspective is from a business view, rather than from a technical or other view.

7.3.1. Identifying Business Processes and the Impact of Loss

The first step in risk analysis is identifying and prioritizing which business applications and processes are needed to run the business. This makes it easier to define the loss if, for example, a businessperson is asked to quantify an invoicing process or a goods receipt process. It is harder for a businessperson to define the loss of a computer. In addition, doing this from a business view clarifies potential losses and the amount of money required to prevent them.

This step is normally performed through a process of facilitation and questionnaires involving line management of each business function and IT. The result is a list of defined business processes, the losses incurred if the processes are not available, and the timing of the losses. The time factor is important because it affects mitigation. When a business process is lost, the monetary impact does not start immediately. For example, the loss of e-mail has a delayed impact because there are other ways to communicate.

It is helpful to look at all sources of loss to either a business asset or a business process using a loss matrix as illustrated in Table 1-1.

Table 1-1
Sample loss matrix

Asset or Process	Financial Loss	User Disruption	Legal Impact	Confidentiality	Embarrassment
Disclosure					
Modification					
Unavailability					
Destruction					

For an impact analysis, this matrix can either contain quantitatively measured impacts that project financial values of the loss, or qualitative assessments (e.g., high, medium, or low) that approximate impacts into general categories. The use of such a matrix ensures that all of the potential sources of loss are identified. It is easy to forget, for example, that embarrassment due to data modification can cause a real loss to a business.

7.3.2. Identifying Critical Assets

Business management should define the criticality of each business process or application. How long can the business run without the process and, as a direct consequence, without the asset for the process? This ensures that business management drives the requirements and that IT management provides the input for the assets needed.

During this step, the assets used to run each business process are identified. Typically these are IT assets (e.g., servers and networks), but they can also be people, processes, and tangible business assets such as buildings and data centers. Certain assets may be required for multiple processes. The intermediate result here is identification of business criticality - either on a subjective scale (e.g., high, medium, low, or scored) or an objective scale (e.g., $1,000 USD loss per hour). Objective figures may be available if a business manager can determine the monetary loss when a process is not available.

7.3.3. Identifying Threats to Assets

For each asset in each business process, the next step is to identify the potential threats to the asset and how often (in probability) they may occur. For ease, it is common to rank the priority of each business asset and deal with the most critical first. A major weakness in this area is that if objective probability data is not available, quantitative methods cannot be used. Qualitative methods should be used instead.

Table 1-2
Example risk analysis for invoicing system

Business Process	Impact of Loss	Impact Starts	Assets	Threats
Invoicing system	High. $30,000 per day	After 2 days	Data center	Fire in the computer room
			Network infrastructure	Network intrusion
			Windows mail servers Directory servers Application servers	System failure
			Operations staff	Critical staff missing

7.3.4. Result

The end result of the risk analysis is a list of business processes, corresponding assets, and the threats to each asset. Table 1-2 gives an example for an invoicing system. The key point here is agreement between business management and IT.

7.4. Security Control Architecture

This step attempts to determine the residual risk, which is the remaining risk after a mitigation plan is applied, for each threat. There are multiple ways of lowering the impact of a threat, ranging from avoidance to a mitigation plan. Mitigation plans can be based on process or technology. Within the process and technology categories, there are multiple mitigation options with differing effects and levels of residual risk. Table 1-3 and Table 1-4 outline two examples of threats and the level of residual risk corresponding to the mitigation options for each.

Table 1-3
Example #1: Threats and residual risk after mitigation

Threat	Probability	Business Impact	Mitigation Options	Residual Risk and Cost	Controls
Intrusion into the network and loss of data confidentiality	Generally high	Embarrassment, legal action if data is private, governance impact-over all high impact	Firewall at the perimeter	• Low cost but high residual risk • Requires regular updating and the use of good processes	
			Encryption of all data	• High cost, low residual risk • Some application	Security policy on data use, labeling of data sources
			Firewalls on all machines	• High cost, medium residual risk	
			No access in or out permitted	• Low cost, high business impact	
			Network partitioning	• Medium cost, medium residual risk • Requires new management tools	

Table 1-4
Example #2: Threats and residual risk after mitigation

Threat	Probability	Business Impact	Mitigation Options	Residual Risk and Cost	Controls
Invoicing systems unavailable due to denial of service or failure	Medium	Very high Direct losses sustained due to inability to invoice Impacts accounts receivable and direct sales opportunities	Manual invoicing system	Low residual risk but high cost of implementing and training	
			Hot swap system	Medium cost but medium residual risk since it covers systems only	
			BCP including all related systems	High cost but zero residual risk and alignment with governance requirements	BCP plan rehearsals, successful completion of training for all staff

In these examples, there are multiple choices of mitigation and residual risk. The key point of using this methodology is that business managers base their choices on their definition of acceptable residual risk. This is stressed because the choices are often technical decisions.

The output from this step is a set of mitigation choices per threat, an approved set of residual risks, and a set of controls that determine if the mitigation plans are working. In the examples given in Tables 1-3 and 1-4 (see highlights), business management chooses:

- Example 1. Encryption of all data: High cost, low residual risk, and some application impacts.
- Example 2. BCP including all related systems: High cost but very low residual risk and alignment with governance requirements.

Ideally, for each threat, the business impact, the risk mitigation costs, and the residual risk position should be stated in monetary terms.

7.4.1. Controls

When choosing controls, it is common for organizations to have too few and to audit these only. The critical issue is to quantify the impact of a control - either quantitatively (which is difficult due to a lack of data) or qualitatively (which is easier but imprecise). Ideally, each control option has an associated cost and impact. The choice made by business management is based on the cost of the control versus the loss (with no control in place) and the acceptance of the residual risk.

There are many dimensions to controls. Missing just one can mean that the overall control system will not work and security governance will fail. As an example, Figure 1-4 shows a variety of controls and three examples of their use.

In the case of the physical example, consider the controls necessary to stop someone from entering a building. Normally there is a sign, a door, a lock on the door, an alarm if it is opened, guards behind the door, guards within the building, and if anyone is caught, some kind of sanction. The same set of control types can (and should) be applied to all situations. The remaining examples show how to do this for processes and technical issues.

Figure 1-4
Types of control and use example

Control Type	Process Example	Technical Example	Physical Example
Directive	Checklist	Banner	Sign
Preventive	Separation of Duties	Authentication	Lock
Detective	Checksums	Logging	Alarm
Corrective	Audit/Rollback	Interrupt	Guard
Recovery	Audit	Rollback	Multiple Guards
Deterrent	Job loss	Disconnect	Sanction

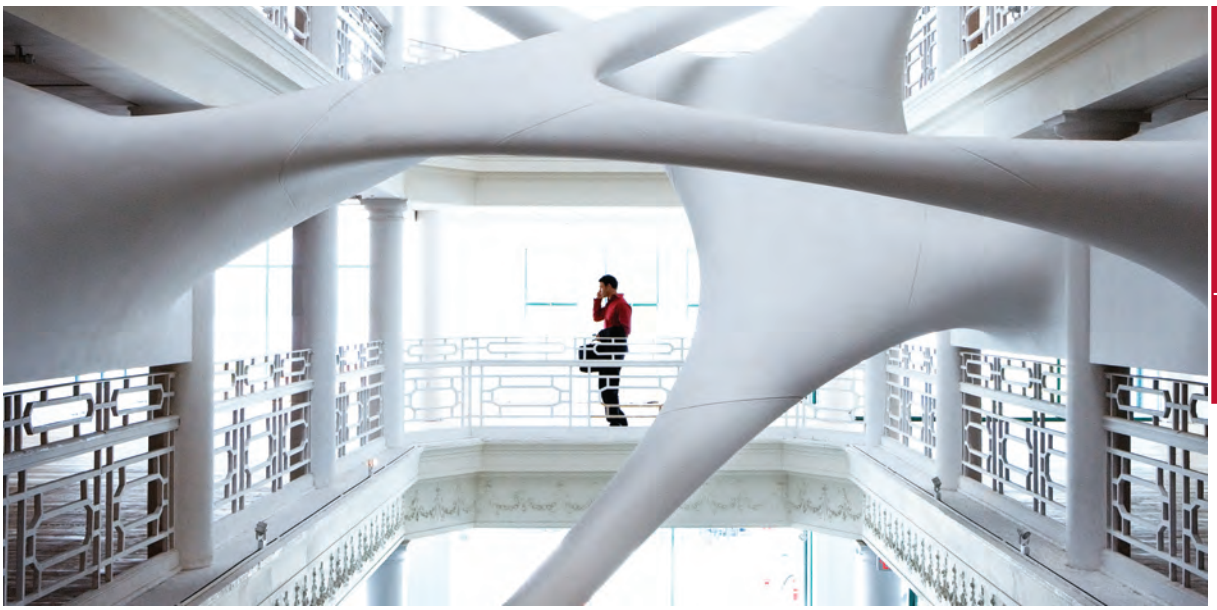

IT is part of the general concept of business control, because IT exists to implement it. As an example, prior to accounting systems, a business transaction was made in cash or in kind. Control was exercised immediately and the sanction was often severe - especially if a promise was made and not kept. Later, accounting systems were introduced to record transactions. Auditing these transactions on behalf of owners or shareholders followed. However, the basic concept of an invoice to help control the transaction and to ensure payment remained throughout. The advent of invoicing faster and more efficiently by IT has not changed the underlying purpose, which is to control loss. (No invoicing probably means a loss.) Therefore, the IT-enabled method of invoicing can and should be cost justified in terms of loss.

The same logic applies to e-mail. E-mail helps make business work and implements control. As an example, if an invoice is in error we can use e-mail to communicate. Although it is hard to justify not having e-mail, a different view is necessary when it comes to determining business risks. A business impact analysis of e-mail generally reveals that it is not critical to business because there are other ways to communicate for business purposes. After some time, the lack of e-mail may radically affect a business, but this rarely occurs in the short term. This emphasizes the concept that analyses should always consider the business impact over time.

7.5. Security Implementation Architecture

The previous step defines the parameters for a security architecture. The outputs can be technical components, topologies, technical choices, training plans, continuity plans, security policies, process definitions, and job descriptions. All of the outputs must be balanced and implemented to ensure security governance works as planned.

Business management should make all of the choices up to this point. However, the implementation step of the lifecycle is a collaboration between operational and IT staff to define the architecture that supports the business choices. The architecture should also produce the controls required for tracking.

At this stage, compromises that require feedback to business management may be necessary, and unforeseen issues may arise that require business decisions. Similarly, some choices will annul others. As an example, a firewall that does not have administration processes defined will soon become useless; the processes and associated costs should be identified. A continuity plan that does not have a process for rehearsing and maintaining it will also become useless. Additional costs identified here can invalidate previous decisions.

This step examines the requirements, defines the overall architecture required, and produces a skeleton or high-level plan. It is important to note that the decisions driving this are business-related, not technical, and that the architecture is not fixed in time. It should be designed with an eye to the future. Within this step, it is very important to avoid designing an architecture that is self-defeating in terms of agility and future requirements. For these reasons, defining and following key performance and goal indicators is critical.

Due diligence is the objective, governance is the process, and compliance shows that the process was followed.

7.6. Implementation

Operational staff and IT staff drive this step, which takes the output from the previous step and implements it. This is usually a team effort and a multidisciplinary team should drive it. Operations people are principally involved in process design and implementation together with training plan implementation. IT staff implement the approved technical architecture.

As always, additional discoveries may require changes to previous decisions. Under no circumstances should changes be made without business management input. For example, it is common to find residual risk changes at this point. If this occurs, the changes should be presented to business management for a determination of the appropriate course of action. This loop is necessary to align business and IT actions - an essential (and often very difficult) part of governance.

7.7. Support, Manage, and Operate

After the overall implementation plan is complete, it is operated as planned. The function here is to execute as agreed. This includes handling events and incidents, managing efficiently, and reporting any issues that could improve efficiency and performance. The objective is not to execute this function as though it were a business with decision - making power of its own. Unfortunately, this improper form of operation frequently occurs.

7.8. Audit and Test

This step is actually developed during the previous steps. The ability to test and audit should be identified as one of the control choices. As an example, intrusion detection capability requires technical choices in terms of network topology, components, and machine agents. It is a separate control step because regular action is required. Everything should be tested for intrusion and vulnerability at least once a year, and audit requirements are often shorter. Common sense dictates a quarterly cycle. Compliance requirements differ based on the industry. In general, KPIs will have been chosen and should be monitored or generally available. In a fully agile company, KPIs should be available in close to real time. Given that compliance drives the requirements in many cases, there should be a continuous effort to produce performance reports that support the needs of the approved compliance framework.

Most compliance frameworks have a high degree of subjectivity and tend to be based on best practices and management judgment. For practical purposes, there are COSO, COBIT, and ISO 2700x. There are also some specific guidelines per industry, but the only one with some degree of prescription is the Basel II Accord, which applies to the finance industry and provides recommendations on operational risk and clear prescription of the need for continuity plans.

7.9. Review and Update

The security and governance framework should be formally reviewed and the results documented on a regular basis. A formal review should also occur for a major change in business structure or a new business venture. The documentation demonstrates two objectives: management approval and due diligence to regulators. For this reason, the process should be formal and regular.

The process of review usually includes the risk position and the results of KPIs, audits, and tests. The results should be reviewed at this stage - not delayed until the next review cycle. Actions related to the results of audits, KPIs, and tests should occur as quickly as possible, depending on the nature and severity of the results.

8. Managing Governance in Practice - Information Security Service Management

Given the plethora of regulations and the general tendency to provide guidelines or recommendations rather than prescriptive advice, how can a business implement a corporate governance framework which is easy to manage and fits their requirements? There is a tendency to rely upon achievement of some standard (e.g., COBIT, ISO 2700x) or to demonstrate compliance to the requirements of a regulation (such as Sarbanes-Oxley). In the best case, both of these methods provide only a snapshot of a business at a point in time. They do not provide an operational management framework for continual compliance, nor do they naturally align with any standard management tools. This makes governance programs potentially dangerous and expensive in practice since there is a tendency to ignore the underlying problems once the standard has been achieved.

HP's Information Security Service Management (ISSM) process is a comprehensive approach to designing and deploying an enterprise Information Security (INFOSec) program with a focus on business impact, control systems design, metrics reporting and operational alignment. The ISSM Reference Model is the culmination of years of HP security consulting experience gained through the design and deployment of hundreds of security engagements. ISSM forms the foundation of any corporate governance program and as such, directly defines controls in line with business risk and provides a method of tracking how these controls work in everyday business operations. A yearly review of the business impact and risk tolerance position combined with control refinements and continuous and stringent everyday reporting provide the key elements of governance.

ISSM is a process-focused discipline for business controls that prescribes capability maturity levels, governance, KPIs, enabling technologies, and service management, and is based on a set of fundamental requirements for any security program:

- Security should be standards-based
- Security is a shared service
- Security is deeply integrated within IT
- Security controls should be highly structured
- Security should be managed through KPIs
- Security reduces operational risk
- Security should be a continuously improving process
- Security should be linked to operations
- Controls should be tracked by KPIs

Aligning controls with accepted risk is a fundamental of all governance and compliance requirements. Two key concepts of the ISSM approach are the use of a maturity model to look at the current and desired future states and the basic use of the ISO frameworks as a guide but with one important difference. When a normal ISO 2700x assessment is done, the controls are looked at as being true or not true. In reality, a control can exist but not be well managed, so broadening the assessment criteria is essential to knowing whether the control will actually work in practice. ISSM adopts the following model for assessment:

- P1: People: Assigned staff to oversee and manage controls
- P2: Policies and Procedures: Governance documentation used to specify and manage controls
- P3: Processes: Operational sequence of activities or events designed to reduce risk
- P4: Products: Defense-in-depth technologies/solutions used to manage or mitigate risk
- P5: Proof: Metrics or validation methods used to track control effectiveness

Only when looking at these five dimensions can one be assured that a control is fit for a particular purpose and that it actually works. Combining this approach with maturity models provides a sound operational base for the implementation of operational governance of security and operational governance of business controls.

Figure1-5
ISSM process cycle

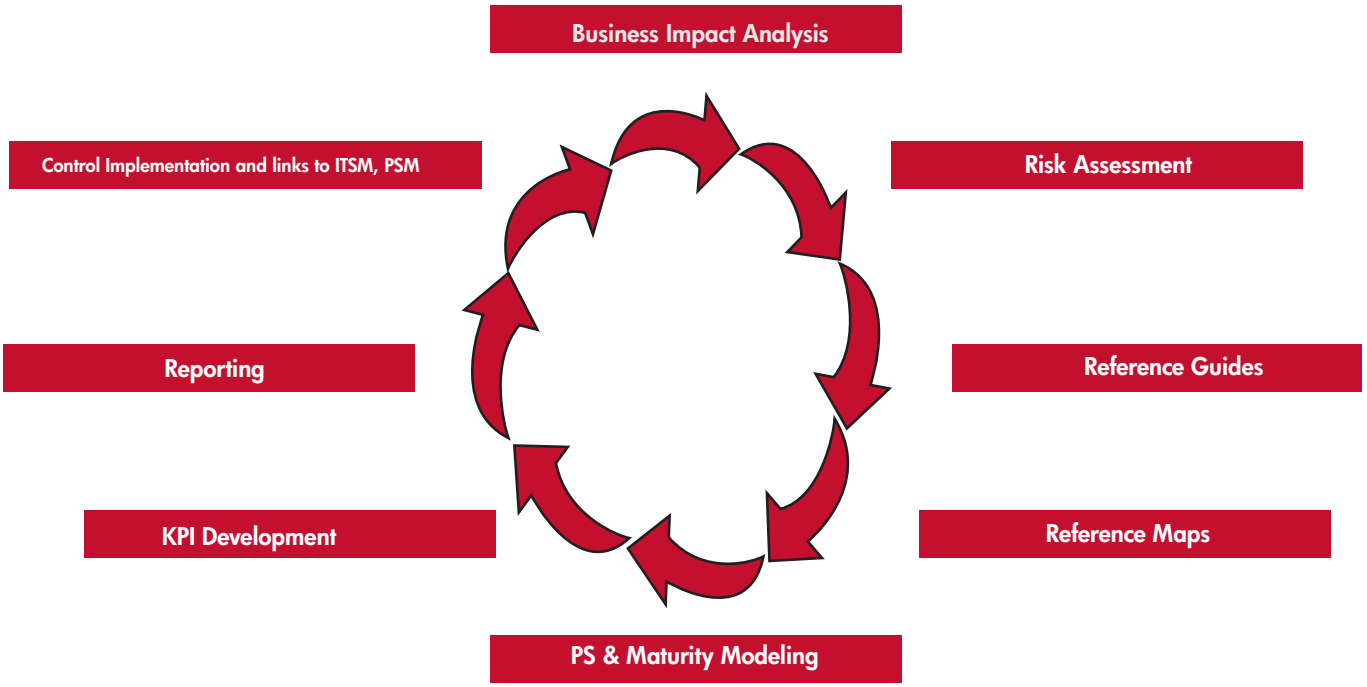

8.1. ISSM Control Model

Figure 1-5 shows the overall process cycle used in ISSM. A business impact analysis shows which areas are of concern and a risk assessment is done for these. Controls are then defined based on best practice from the reference guides and in line with the particular regulatory framework from the reference maps. The current and future view of implementation is defined according to a maturity model and the P5 criteria (above) followed by the development of KPIs for those areas needing tracking. Linking these KPIs to operational data sources completes the cycle.

The following major domains are addressed within ISSM:

Business Impact Analysis (BIA)

Identifies the critical business functions and services and quantifies the impact of their loss to the business. The output from this stage is a BIA which gives a broad picture of how much should be spent to safeguard against risks.

Risk Assessment

Without quantifying and knowing the risk in financial terms, a business fails in four areas - it does not fulfill any regulatory or due diligence requirements, it does not know whether to accept or reject the risk, it does not know the potential impact on its business, and it does not know how much to spend to mitigate the risk to an acceptable level. This step should be formal and for due diligence purposes, done yearly and signed and agreed by management.

The output from this stage is a risk assessment which simply shows the threats, vulnerabilities and risks a business faces and which complements the BIA to provide a complete picture of how much should be spent but not where. A key element here is that the actual risk appetite of a business needs to be taken into account since some risks can be accepted simply because their impact is low (or not material). Not all risks need to be mitigated and not all need to be tracked.

Security Control Framework Definition
This ensures that all domains are covered and that we use a standard model for defining security domains in which controls should exist - for example; asset management, physical security, etc. This step is essential to avoid the situation where an information security program only covers some of the domains affecting security or only covers IT, and as such misses the holistic approach to controls that is needed. As a simple example, if an information security program focuses on access control but does not look at human resources, access could be granted to people under notice or staff could be employed who represent a risk. Making sure all domains are covered and work together is the most difficult part of any security program and not taking this into consideration is the cause of most failures.

For each domain, the underlying disciplines and control elements are reviewed, drawing heavily on the structure of the ISO 2700x security standards. Predefined reference guides take into account other standards, legislative considerations and best practices.

To understand the required underlying controls, a cross reference which shows which ISO elements apply to which framework is required. ISSM includes a set of guides mapping well-known industry frameworks to ISO elements. The purpose here is to ensure that all areas of the target regulation are covered. These authoritative guides provide detailed deployment knowledge for implementing compensating controls. Reference guides are selected from leading industry and security standards bodies, making them universally accepted by all auditing and regulatory entities. Each guide has been carefully researched for applicability to each compensating control and provides the depth of knowledge to be used in whole or part depending on the desired maturity level of a compensating control.

The output from this stage is a control framework which shows which areas need to have controls to mitigate the known risks.

Control Definition
For each ISO discipline there are many types of controls that could be put in place and these are specific to the business and the risk. To help in implementation HP uses pre-defined control implementation templates that show the typical controls that can be used for each discipline within each domain. It is very common for the controls applied to risks to be specific to IT or to a particular technical domain. It is critical, however, that there always be multiple controls for each risk and that these controls be a combination of technical, managerial, process and accounting controls.

The control guides suggest ranges of controls for each discipline. As a general comment, control should be chosen from the families and there should always be more than one per risk:

- Directive: An order such as “private system”
- Preventive: For example, access control
- Detective: Such as an audit log or a checksum
- Corrective: Such as a rollback mechanism
- Recovery: Such as an audit system allowing recovery
- Deterrent: Such as a sanction

The outputs from this stage are a Statement of Applicability and the mitigating controls required for the business risks and regulatory requirements.

Measurement and Reporting

With the foundations laid from the above steps, risks are reviewed, control areas are agreed and comparisons made with best practices in terms of controls and regulatory requirements. Now an assessment and gap analysis can be executed which will show the current state and the desired future state in terms of control effectiveness and maturity. During the assessment phase, KPIs should be developed to track the implementation and performance of controls. This task is largely a question of management judgment. KPIs can be split into two groups - direct and indirect.

A direct KPI could be the number of false login attempts being below a certain threshold. An indirect KPI could be measuring the uptime of a system as a means on inferring that the controls running on that system are working and therefore that the control works. This would be relevant to measure an application control, for example. Some KPIs could be manual and manually tracked and manually entered (such as the number of people checked into a building).

Management should choose KPIs based upon the criticality of the control and the available data to track it. It is important to recognize that KPIs represent tracking of the controls to mitigate business risks and that the number that can be sensibly tracked should be limited to the number of business risks which are material to an enterprise. Having a KPI for every risk and tracking every control in the ISO 2700x standard can be a waste of effort if only because the vast majority are not material.

Important Note: In this context, "material" means that the impact would not have a materially significant impact on the financial position of a company.

The output from this stage is the set of KPIs which need to be tracked.

Link to Operations

In a typical control framework there is often a missing link between sources of operational information and the control system. This happens when the control system is developed independently of the operations side of a business. Typically, rich sources of operational data which are relevant to controls and risk management can be found in areas such as change management, continuity management, availability management etc. ISSM provides the link to these areas using ITSM/ITIL v3 and provides an effective operational implementation of risk control providing many benefits in terms of cost reduction and effectiveness.

As an example, a mitigating control could be an application running on a specific resource. If the resource is not available for a time period higher than that planned, it can be deduced that the control is not working and therefore the business risk is not mitigated.

Other sources of control implementation and tracking come from a security management system (covered in detail in Chapter 2 of this book). A security management system should normally consolidate all of the information from the technology and processes managing security and synchronize actions between all of the functions in a company. In addition, a security management system should proactively act to avoid security problems before they impact a company rather than just managing the damage they cause. The HP Proactive Security Management framework is a good example of this.

Lastly, some controls will need to be tracked manually. For example, for physical access it may be relevant to use entry logs as a KPI.

Linking the control cycle above to operations ensures that information feeds from existing or future business operations will be used to keep the control system up to date. For example, change management should be synchronized with the control system to ensure that changes do not break the controls in place.

It is important to remember that while mitigating controls should be in place for all areas of risk, it is neither feasible nor necessary to link them all to operations. Those linked to operations should be those linked to mitigating the highest business risks and impacts (those which will have the greatest material effect on financial results). Only these require continuous management attention.

The output from this stage is an Operational Security Management framework where the linkages are made to service management systems and security management systems.

9. Moving to Continuous Compliance

Traditionally, auditors review a control system on a regular basis, and a yearly report shows that financial results are correct. Information is also produced to demonstrate regulatory compliance. Relatively recently, however, compliance regulations have required more objective demonstrations of control systems. In addition, Sarbanes-Oxley requires declarations of control system failures, with explicit management approval of the existing control system rather than auditor approval.

Not all organizations and control systems are directly affected by the declarative nature of Sarbanes-Oxley. Most organizations operate in jurisdictions where control system failure is treated in a much more relaxed manner. This will likely change regardless of an organization's legislative framework, however. There is worldwide pressure (provoked by Sarbanes-Oxley) to conform to a more objective standard.

An opportunity now exists, under the guise of regulatory requirements, to redesign control systems to provide good control and compliance at a lower cost and with early warning of failures and deviations. To accomplish this, control systems must move from traditional, standard compliance methods to continuous, real-time assurance.

9.1. Comparison of Standard and Continuous Compliance

Most control systems tend to be static. For example, in a typical system:

- A process is designed to produce reliable results for financial declarations.
- The process and its implementation components are audited to see that they work.
- Checklists and statistical methods help to monitor the control systems.
- Results are gathered (often yearly) to enable a declaration that the control systems are doing their job and that the financial results produced are correct.
- Process failures either occur as events during a financial year or are noticed at audit and corrected then.
- Process failures tend to be due to people failure, and the results are often dramatic (with some exceptions). Process redesign is rare.
- An unspoken rule dictates questioning of the implementation but not the control system itself.

By comparison, the primary objectives of continuous compliance are to align risk with mitigation and to provide early warning of risk and mitigation failure. In addition, instead of continually running through checklists or performing audits, early warning of failure is derived from indicators that are assigned to the control system. Table 1-5 compares characteristics of standard compliance and continuous compliance.

Table 1-5
Comparison of standard and continuous compliance

Standard Compliance	Continuous Compliance
• Static, cyclical reviews	• Ongoing assurance
• Historical-based	• Strategic
• Intrusive	• Collaborative
• Point-in-time retrospective	• Real-time transparency
• Unexpected fluctuations in the control environment	• Sustained, adaptive governance
• Coverage-based	• Risk-based
• Adherence to rules	• Response to risk

9.2. Continuous Compliance Example

Examining a typical business function, such as accounts receivable, can help to illustrate the alignment of risk with early warning and mitigation using control system indicators. Accounts receivable is a set of financial processes. The processes use applications, and the applications run on machines in a business network. Performance indicators can be assigned to each level of the accounts receivable function:

• Financial process. An inappropriate document type such as a cash receipt for a debit serves as a financial process performance indicator.
• Application. Modification of customer credit terms is an application performance indicator.
• Infrastructure. Uninstalled security patches and lack of system availability are infrastructure performance indicators.

Determining whether the financial process uses inappropriate document types is a direct measurement of whether the process is working properly. For example, a cash receipt is not normally issued for a debit, indicating that something is wrong. Monitoring whether customer credit terms are modified with an application is a direct measurement of financial impact. For accounts receivable, credit term modification is usually a sign of client weakness. Frequent extension of credit terms reduces the reliability of the financial declaration of accounts receivable. Note that we do not need to review the entire accounts receivable system to discover certain events that can affect financial results. This is an important distinction. The application or the application process can operate correctly, while the financial system operates incorrectly. This is a critically important concept; it is the point at which we cross from security governance to compliance. In the example, the third level of measurement is infrastructure performance. If we measure system availability, we have an indirect measure or indicator of the health of the underlying infrastructure, which is part of this overall business process. System availability can be measured automatically and often.

9.3. The Efficiency of Continuous Compliance

Generating performance indicators for measurement assumes that general best practices are in place. In addition, it changes the audit and compliance process from a point-in-time audit to a continuous process method that provides early warning of trends, risks, and corrective action. Checklists could be used for compliance, but this presupposes that the checklists are continuously updated and are looking for trends. Use of checklists also presupposes that the underlying control system mitigates the business risk fully and correctly, and that it can provide compliance data.

Compared to checklists or other static methods, monitoring continuously is much more efficient. This requires input from multiple sources:

• Direct input. Direct input from applications usually requires application changes. Most applications have simple checks like range checks or exceptions (such as invalid customer account numbers). This data is presented at application runtime, but it also needs to be presented to a system that correlates all control checks.
• Manual input. This type of input consists of manual process information such as results from audits.
• Indirect input. System availability is a good example of indirect input. If a system in a control infrastructure is running as expected, one can assume that the controls are also running as expected. If the system is monitored and unavailable, it is an indicator of potential control failure. Another example of a system that can be monitored automatically is an intrusion detection system.

In a normal process for continuous compliance, the following items are identified:

• The acceptable level of error
• The financial statements (sub accounts) that could produce the biggest errors
• The applications that produce these sub accounts.
• The infrastructure (in the broadest sense) that runs the applications and processes
• The controls currently in place
• Whether the controls are designed and working correctly to catch errors
• How the controls are reported and correlated.
• The indicators that lead to trends requiring actions to avoid misstatement

In compliance, we look at a control system to determine the compliance impact. This equates to establishing which controls, if lost, would result in a material misstatement of results. The overall mitigation cost should be less than or equal to the level of material misstatement. In general terms, the materiality level is quite high. In the majority of corporations (millions of dollars). This means that fairly large errors can occur before the governance and control system experiences a failure which would impact compliance. The investment in mitigation methods (IT controls, general controls etc) should be proportional to the potential loss and also proportional to the potential misstatement of results. For example, it could be that an asset is worth ten thousand dollars and its loss could cause fifty thousand dollars of damage - at first glance it would appear that the mitigation cost could be fairly high if the asset could be easily damaged regularly. If there is a 30% chance of loss every year then the mitigation cost of fifteen thousand dollars is justified. However, if the materiality limit of the company is five hundred thousand dollars and this asset loss would never cause such an impact, it may not be worth protecting the asset at all. The decision here is a management judgment based on risk appetite.

10. Using Models and Model-based Technologies to Support Security Governance

Traditionally, best practice for assurance management employs a risk-associated "control" architecture for the IT environment. An associated lifecycle directs testing and adapting those controls as the environment is operated and changed. Problems are monitored within the contexts of the control architecture and the efficacy and interdependency of the controls. However, this is extremely difficult to achieve in an ever-changing, complex environment.

A more progressive way to deal with assurance management is to model the control framework. This immediately lifts the assurance lifecycle from a series of people-based processes (risk management, control design and implementation, audit and review) to one where model-based technology enhances, connects, and (where appropriate) automates the process. This brings the benefits of efficiency, consistency, improved communication, and ultimately more control and assurance.

One of the major issues in achieving this state is the fact that control systems exist in many forms. Inconsistencies are common, and some parts of the true control system are not considered. For example, physical access control to a building is in theory part of the financial accounting audit, but it is often omitted in practice.

To address the challenge of having many forms of control systems, a model-based assurance framework can provide new information via a high-level view. Modeling provides experts (such as auditors, assessors, and security architects) with an indication of completeness, a more thorough view of the interaction of controls, and more easily managed control indicators. For example, to mitigate the risk of people inappropriately accessing applications, a process control ensures that each addition of a user or role is approved and fits with security policies such as segregation of duty. The model captures the relevance of this control to the overall control framework and the correct operation of the process, resulting in appropriate reporting of exceptions and metrics.

HP Labs has developed a toolset to support a model-based approach. The components of the toolset include:

- A model development tool that allows security specialists, auditors, and/or risk officers to graphically build up a model of the control framework.
- An analysis engine that operates on top of an audit database where all of the specified information is gathered. The engine and model drive analysis that shows which controls (according to the model) are working effectively and which are not.
- A reporting engine that presents the results of the analysis as a navigable dashboard that highlights areas in the control architecture requiring attention.

A modeling approach can provide a continual view of risk or overall compliance. It can also automate much of the auditors' fieldwork. This supports continual compliance much more significantly than merely checking whether technology is configured or working properly. It actually reports against controls that directly relate to risk, automating much of the people-based assurance loop.

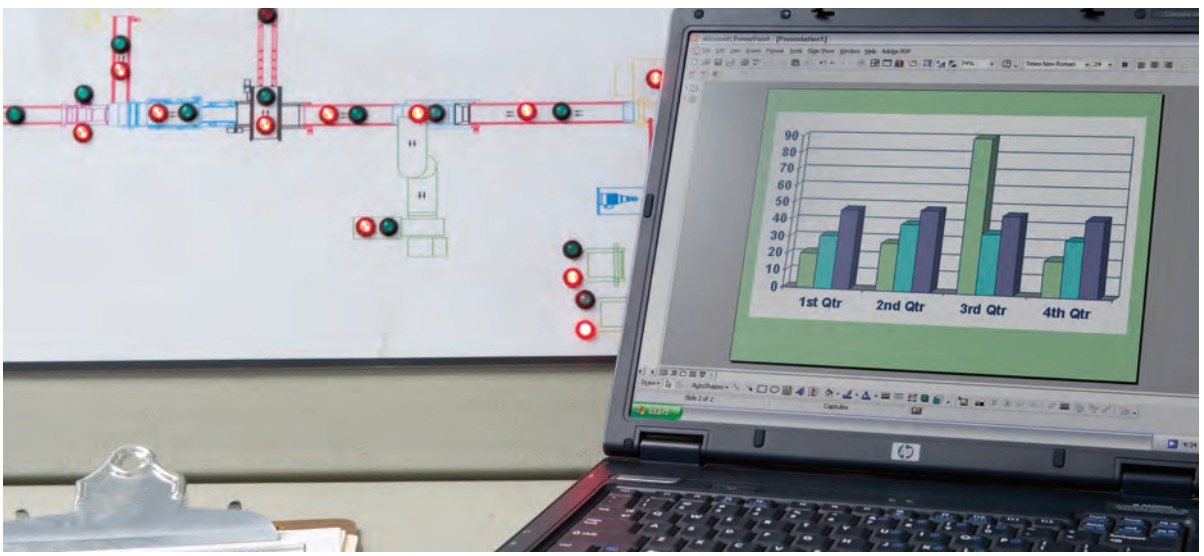

11. The Economics of Security: An Example

In theory, any money spent on security should be directly related to the risk. For a given risk and therefore, a potential loss, the cost of its mitigation should not be more than the potential loss over time or the value of the asset at risk. A suitable timeframe for risk determination is three years. Mitigation never eliminates risk entirely and any remaining residual risk should be of a level which is within the risk appetite of the business. However, when calculating the potential for losses, the losses should incorporate the complete impact on the business and not just the damage that is immediately visible.

How Eonomical is Security?

There are two non-avoidable problems in managing risk. First, a business has the classic problem of not wanting to take too many risks and being unable or unwilling to express their desire to take on risk in terms of a risk appetite. Second, we have the problem of getting accurate and measurable data on the probability that an event will cause a risk to materialize.

Even if a business can define its risk appetite in terms of business priorities, there remains the thorny problem of the residual risk implied by each of the mitigation options. There will always be one or more ways to reduce a risk to an acceptable level and each approach has, or should have, a different residual risk.

The economics of security can be illustrated in the way that companies protect their PCs by securing them to tables with cables. A cost/benefit analysis shows this approach might not be justified.

Let's say the PC costs $1,000 USD. (The figures used in this example are examples but the method applies to all security decision making.) There are 500 PCs in the office and three are stolen or lost per year either from the office or when people travel. This means the annual loss expectancy is $3,000. However, this is not the full story because this figure includes only the cost to replace the hardware. To build a complete cost picture, there are also other costs to look at:

• Inconvenience: The person whose PC is lost or stolen now has no PC and must either borrow one or order one to be able to continue working. In either case, there is lost productivity that can be calculated by assigning a monetary value to each hour of lost time.

• Recovery costs: When the new PC is delivered it needs to be configured with an operating system and applications and the user has to recover its data, assuming that the data is available in a backup. Again, the cost of the time required to configure the PC to a point where the user can become produ tive can be expressed in terms of a number of hours multiplied by the cost of the user's time.

• Loss of working time: The person either cannot work or works less efficiently. Each user is different and the cost depends on the tasks that the user performs and their level within the organization. The loss of a laptop used by an executive who is currently negotiating a critical deal is clearly more expensive than having to replace an older desktop that is used by a more junior member of staff.

• Loss of data: This could be fairly expensive either because no backups were made or because the data is worth something. The loss of confidential data may impact a company's reputation, especially if that data relates to customers.

• Reputation: If loss of a PC and data happens regularly, the company's reputation will probably suffer.

• Liability: If reasonable actions for protection are not taken, there are potential liability costs. For example, if the data loss concerns personal identifiable information (PII) or if data protection is required by government or industry regulations.

The latter two costs can be very high and are very difficult to quantify.

What mitigation options are there in this scenario? The threat is mainly one of theft or accidental loss resulting in the physical non-availability of the device. In our scenario, there are two options - a cable to attach the PC (not much use in a taxi and of limited use in a hotel or airport) and/or encryption software to protect the data if it is lost or stolen (since data loss represents the highest potential loss either to reputation or due to the data value). There aren't really any other options to mitigate the risks and to reduce the losses due to theft. A backup solution is also needed to recover data in the event of loss.

Cables to secure PCs cost only $20 but their real cost is much higher. If cables are purchased externally and then distributed to users, there is process cost to manage the purchase and distribution. If users order cables personally then there are hidden process costs such as the time required to complete, submit, and authorize the expense claims for the cables. Both these processes are actually expensive in terms of the time required to order, install, and manage the cables and it would not be unreasonable to use a figure of $300 per cable since this represents a typical level of accounting control (or cost per transaction).

Table 1-6
The economics of security example: summary of costs (U.S. Dollars)

Loss item	Yearly cost of loss	Lifetime cost	Mitigation method	Mitigation cost total
PC hardware	$3,000	$12,000	Cable	$150,000
Data	$30,000 (3)	$90,000	Encryption	$200,000
Reputation	$30,000 (3)	$90,000	Communications plan	$200,000
Recovery	$2,700 (1)	$10,800	Backup	$100,000
Inconvenience and loss of time	$18,000 (2)	$72,000	Cable/backup	Included above
Liability	?	?	All	
Totals		**$274,000**		**$650,000**

Encryption software costs roughly $80 per license per PC in volume. However, the real costs here are higher too. There are costs associated with:

- Software purchasing and renewal
- Software configuration
- Support for users
- Training for users
- Key management

Even if it only takes a user one hour to load and understand how the encryption software works, the average cost to deploy the software is $150 for a professional. Taking everything together, it is not unreasonable to assume a total cost per PC of $400.

For backup, there needs to be some kind of centralized system or a personal backup solution. A reasonable estimate for this could be $200 per user but can be far higher depending on the software and hardware components used in the backup solution. It does not matter if this is the cost and time for the user to deploy their own backup solution or it is the cost to deploy centralized and automated systems plus the necessary storage and administration for all users to access.

For a total of 500 PCs, we have the following costs:

- Cable cost per PC: $300
- Encryption cost per PC: $400
- Backup cost per PC: $200

Let us also suppose the solutions work for four years for a single one-time mitigation cost. The costs are per user and needed for all 500 machines. To estimate recovery costs, we can assume that in practice, recovery takes 6 hours and that the loss of a PC usually leads to a week of inconvenience and lost time. Using an hourly labor rate of $150, we have an inconvenience cost of $6,000 per week per PC lost or $18,000 for the three lost PCs per year.

Let's also assume $10,000 in losses for data and reputation (each).

This brings the total mitigation cost to $650,000 for a total loss expectancy of $275,000 (here we simply use the lifetime cost of the losses in the table above).

The results of this risk analysis would be:

- We are spending too much to mitigate the potential loss.
- The mitigation cost of losses due to theft and physical loss (cables) far exceed any reasonable cost for the potential loss.
- The obvious solution is to deploy encryption since it protects against the biggest losses.

This type of analysis should always be done when looking at risk and mitigation costs and sometimes the answers are surprising. The obvious and visible solution to a problem does not always stand up to close analysis when looked at from a risk point of view.

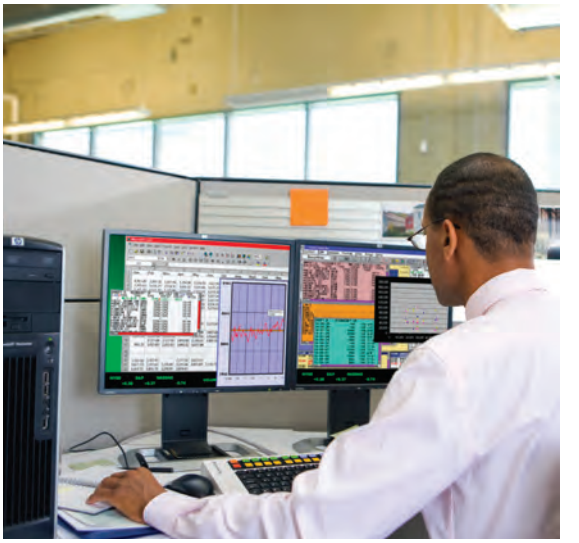

12. Key Performance Indicators and Metrics

Using metrics to measure the effectiveness of mitigation methods can be fraught with danger. It can happen, for example, that too many metrics are tracked, making it impossible to gain a complete picture of effectiveness, or that as a result of this overload of information, the top few metrics are tracked.

The danger here is clear. It is not possible to define the top risks in any sensible way. If the top 50 are tracked, how does one know that number 51 will not be the one that causes unmeasurable damage? Given the trend to zero-day technical attacks and attacks which have not previously been either seen or exploited, it is most likely that new technical attacks will be of a previously unknown type.

If we suppose that we are tracking the top 50, what is a good result? At what point should management attention be raised? Are 48 out of 50 a good result for example?

A metric should more properly be referred to as a Key Performance Indicator or KPI. A KPI represents the performance of a mitigating control or set of mitigating controls which mitigate a risk which is perceived to have a high business impact. Only those KPIs relating to unacceptably high business risks should be tracked for management attention. By definition, there are only a few business risks that need tracking, but the process of deciding which ones places the onus on business management and not on technical management. However, both sets of management need to accept that risk exists and damage occurs, but management attention should be focused on business impact alone. The remaining risks are better handled by simple best practice implementation of mitigating controls.

As an example, the business impact of the loss of an invoicing system is seen to have a significant impact on Days Sales Out (DSO) and on client confidence. For example, the invoicing system runs on three machines, uses four staff, the network, an invoicing application and a room.

We can reduce this by deciding that the staff has coverage from other parts of the company and the room can be anywhere. This reduces the problem in this case (but only in this case since very often the major points of weakness are people and buildings) to one of machines, network and application.

The mitigations for these components could be redundant machines and network, rigorous change control for the application, intrusion detection and anti-virus systems and some audit capability. Whatever the many and varied components used, the general case will be that if any one of these breaks, the invoicing system won't work.
As KPIs for this set, we can choose:

- Unplanned application changes
- Unplanned downtime
- Missed anti-virus scans

Note that these KPIs are not of the type “number of intrusion attempts”. There will always be occurrences like this but they don't necessarily impact anything. The KPIs above, however, are measurable combined events. This is very similar to driving a car. There are always events such as the temperature rising in the radiator. Only at a certain point does this cause an impact. Separating the noise from the impact event is the purpose of choosing a KPI rather than choosing to monitor the noise.

A KPI is entirely dependent on the impact being tracked and should be seen in this light - the key is to map business process to a group of KPIs and to track the group as a group.

13. New Model-based Analysis Approaches to Support Risk Analysis - Trust Economics

Aligning the security architecture which meets the security governance objectives of the business within the overall requirements of IT governance is a challenging task. Another starting point for analysis is based on new research. In HP Labs which aims to address the two key challenges facing CxOs and CISOs with responsibility for information and systems security:

- The poor economic understanding of how to formulate, resource, measure, and value security policies.
- The poor organizational understanding of the attitudes of users to both information and systems security and of their responses to imposed security policies.

The solution is to develop economically and mathematically rigorous systems technologies and tools with which these questions can be addressed.

A rigorous understanding of the behavior of the users of a system (network), together with the economic value of the system's security measures, can be captured within an extension of some established, mathematics-based systems modeling techniques. To this end, Trust Economics is the conceptual framework within which HP is pursuing the study of the economics of information security policies, protocols, and investments. HP's perspective is one of "systems thinking" and, critically, our aim is to seek to integrate the following three perspectives:

- Modeling the behavior of the users of systems, both internal (operators, staff) and external (customers, regulators), in the context of security policies and protocols.
- Mathematical modeling of systems, organizations, and networks, including the security policies and protocols which govern access.
- Economic modeling of the costs and value of security policies and values.

One of the key strengths of this process-driven modeling approach is that it generates *executable* models and thus generates data from simulations that enable hypothetical questions to be asked of the model. Ultimately, the performance and effectiveness of a given security architecture will require exacting approaches to solving the complex trade-off issues arising between relative investment levels in people, processes and technology. For example, modeling patching processes in a business might reveal weaknesses in patch evaluation and deployment.

One solution might be to acquire technology to implement more automated patching; another solution might be to increase or make better use of the people involved in the patching processes. The choice requires careful analysis and the ability to simulate possible solutions prior to committing to a change project.

14. HP Governance Services

Demonstrating security governance is a continuous process. Similar to auditing, most governance frameworks require proof to be presented yearly and in some cases, such as Sarbanes-Oxley, quarterly. HP provides professional services (as well as the necessary hardware and software components when needed) for each step in the governance lifecycle. These services include:

Risk Analysis
- Business impact analysis
- Threat analysis
- Training and facilitation
- Physical security and environmental assessment

Security Control Architecture
- Policy development
- Business continuity planning
- Compliance with control frameworks and legislation
- Mitigation planning

Security Implementation Architecture
- Process design and implementation
- Technical design
- Process design including IT Infrastructure Library (ITIL)
- Training

Implementation
- Systems/networks implementation
- Process development
- Training

Support, Manage, and Operate
- IT Service Management (ITSM)
- Incident management

Audit and Test
- Framework gap analysis
- Vulnerability tests
- Intrusion tests
- Compliance tests
- Audit

For further information about HP Professional Services offerings related to governance, see www.hp.com/go/security.

Figure 1-6
The CIO's balancing act

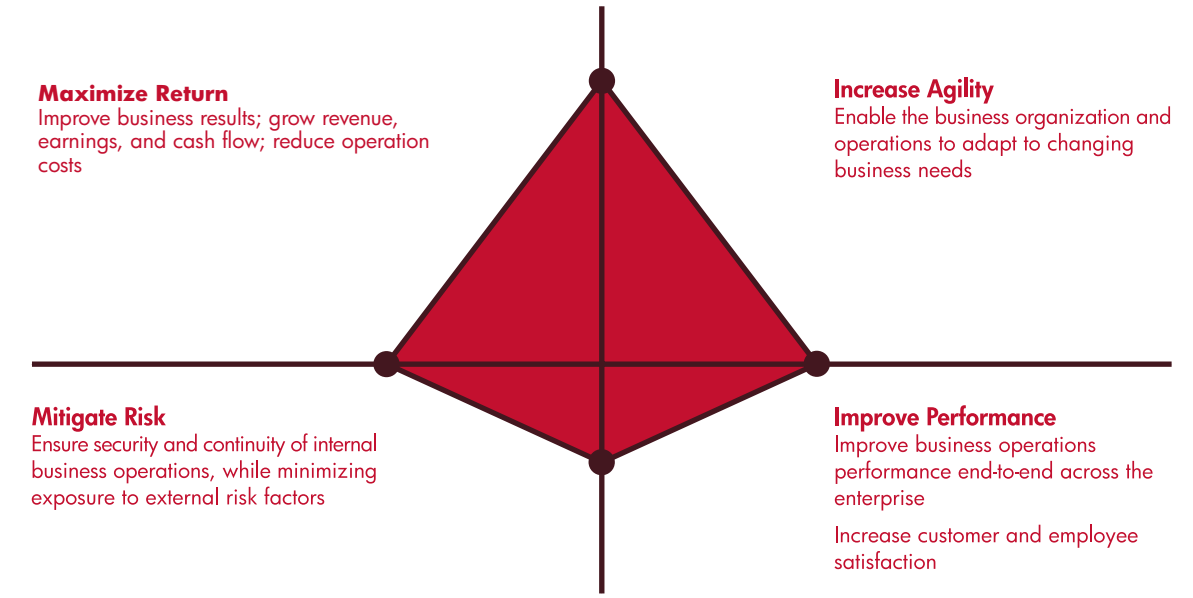

15. Security and HP's Vision

To take advantage of new business opportunities as they arise, corporate governance requires efficient asset management and availability. One of the most striking features of today's business environment is its dynamic nature. Successful companies capitalize on change, turning what is often unexpected and disruptive into a business advantage. The ability to respond to changes can be summarized into four primary imperatives for business and technology: mitigate risk, maximize financial return, improve performance, and increase agility. As shown in Figure 1-6, these four imperatives are interlocked; they simultaneously apply to all governance decisions.

A company cannot gain market agility and operate effectively without governing security as a key business enabler. Conversely, those companies that fail to make the connection between agility and security cannot operate efficiently or react in a timely manner to new business opportunities. This is not a choice. The road to corporate governance and compliance with regulations requires that security governance is demonstrable and that adaptability is a reality.

The roadmap is maintained through the governance lifecycle, based on the objectives issued by the board as well as the analysis of the current and future states.

One of the key elements in the transformational journey is to break down implementation silos. Ad hoc security implementations should be migrated to security services used horizontally across the enterprise and managed in a consistent, secure, and auditable manner. This applies to infrastructure and application control points, for both the service delivery and the service delivery management environments. In particular, the process through which exceptions are managed should be carefully engineered and strictly followed; failure to do so impacts both the ability to adapt and the overall level of protection provided by the control mechanisms in place. This is not just a question of technology; people and processes play a central role in the implementation of the transformational journey.

16. Governance Summary

Business objectives must drive security governance. Likewise, mitigation plans and the costs of such plans should be based on business impact or loss potential. This ensures alignment of costs and time horizons with business needs. Mitigation plans for business risks will always be a combination of technology, people, and process. And the majority of these plans are managed and implemented by IT. Because most mitigation methods rely upon IT technologies and processes, joint teams from business operations and IT should determine the mitigation plans to respond to business risks.

Achieving compliance with regulations calls for good control systems, in general. Building good control systems requires organizations to lay the foundation correctly. Aligning with best practices such as COBIT and ISO 27002 facilitates this process.

HP's Security Governance Services include a broad set of offerings delivered across the governance lifecycle to build an enterprise-wide policy foundation, a secure and agile architecture, process framework, and an organizational structure. Together these services enable businesses to manage the risks associated with their information assets.

Chapter 2
Proactive Security Management

"Security considerations are required across all functions of IT. Only a proactive approach to security can enable an IT function to meet its commitments to protect the IT assets of an enterprise from the many types of security threats in today's technology environment."

-Mike Baker, Vice President, Chief Technology Officer for IT

Proactive Security Management is an important and integral part of an organization's security infrastructure and operations environment, ranging from network-wide controls security to physical security, taking into account both the proactive and reactive aspects of security operations. HP has implemented a Proactive Security Management framework focused on establishing, operating and maintaining security management functions and procedures to support the business objectives while constantly protecting various business services and IT assets throughout their operation. The proactive emphasis for HP's approach to security management also ensures that threats are identified early and the required protection is robust, scalable, and flexible enough to anticipate and adapt to rapidly changing conditions.

In this chapter, we define proactive security management, review the conditions driving its need, and present HP's framework of technologies and services for proactive security management.

Note that while Identity Management is certainly part of the Proactive Security Management topic, HP has chosen to call out Identity Management separately due to its significance and complexity. It is discussed in Chapter 3 of this publication.

1. Definition

The fundamental purpose of security operations management and security products is the protection of business assets. In this security context, protection means providing appropriate confidentiality, integrity, and availability for a set of business assets. Therefore, proactive security management refers to the installation and operation of a set of processes, tools and services to establish and maintain a specified level of confidentiality, integrity, and availability of data, applications, systems, networks, and other IT assets.

To understand the broad scope of proactive security management, there are four parts that make up the whole:

- Managing the protection of data, applications, systems, and networks, both proactively and reactively
- Responding to changes in business and organizational models as well as the changing threat environment
- Integrating with IT infrastructure management and operations
- Maintaining a level of security and operational risk as defined by the organization

1.1. Managing Protection Proactively and Reactively

With the increased complexity of everyday operational security management of IT infrastructure, the growing rate of vulnerability discoveries and the need for regulatory compliance monitoring, it is apparent that reactive security methods, although important, are no longer sufficient. Reactive mechanisms are certainly required in IT environments, but they deal with attacks or other security incidents once they are already in progress - when damage might have already been done and the associated cleanup and lost business costs may already be adding up. Therefore, proactive security management is the natural complement to reactive technologies by providing methods, technologies, and services with the following capabilities:

Reactive Security Management Capabilities:

- Blocking known bad behaviors
- Isolating infected systems
- Enforcing security policies
- Automatically responding to known inappropriate behavior with alarms/alerts for humans and/or predefined, decisive actions
- Performing regular security assessments including manual or automated penetration testing of deployed applications
- Forensic investigations of security events
- Incident management

Proactive Security Management Capabilities:

- Ensuring that everyday IT operations are carried out in accordance with internal security policies and external laws and regulations
- Finding and fixing or mitigating vulnerabilities before they are exploited
- Reviewing web service applications for security defects during their development and deployment
- Automatically responding to suspected inappropriate behavior with cautionary responses to minimize or contain potential damage, then with alarms/alerts for humans
- Modeling solution architectures and networks as well as policies and governance structures to examine proposed changes and their effects on current security and compliance states (and unintended consequences; such modeling is done *before* changes are put in place).

Reactive and proactive technologies cannot provide 100% coverage from vulnerabilities and threats. New vulnerabilities and threats are regularly discovered (and created). Therefore, we strive to optimize the management of the security infrastructure with a combination of reactive and proactive technologies and methods which will build a most appropriate and effective security infrastructure.

1.2. Responding to Changing Business Models

Once a well designed and managed security infrastructure is in place, it must have the ability to adapt to the various threats that emerge and support changes in business models, both internally and externally to the organization. Business model changes can come from organizational changes such as reorganizations or mergers, or from new business opportunities such as new online services. For example, the requirements of proactive security management during a merger might include integrating different security technologies such as intrusion prevention systems and managing employee privilege and authority changes, or a change in operational process may be required. These transitions must happen quickly, taking into account change management, release management, capacity management, as well as security management. Further, it is not hard to imagine that two companies merging would have different policies on authenticating employees, for example, with different authentication products being used.

Proactive security management dictates that the security infrastructure be designed with the potential for such changes in mind, and that security mechanisms are:

- Adaptable to new models to insure compliance and adherence to security policy is maintained
- Extendable to incorporate new security technologies and respond to new threats or classes of vulnerabilities

1.3. Integrating with IT Management

Proactive security management is not an island unto itself - it is a piece of the whole IT and corporate management picture. Therefore, changes in security operational management - whether dictated from a change in business models, laws, regulations, threats or vulnerabilities - must be made to preserve the tight links between security and overall IT management.

For this reason, with HP, proactive security management is tightly linked with the Information Technology Infrastructure Library (ITIL), which is a framework of best practices that promote quality computing services in the information technology sector. Using this ITIL framework, proactive security management can develop all its operating process to manage the security infrastructure and be linked in with the overall IT management structure. Security touches all portions of IT infrastructure and must be integrated from both a technology perspective and management perspective.

Patch management, asset management and configuration management are three examples of IT technologies that have an impact on security. If a critical security patch is not applied before a widespread attack happens the overall security infrastructure fails. But you cannot simply apply any security-related patch without proper patch management techniques to perform testing, and systematic, audited patch application, because that patch might adversely affect performance or integrity or supportability of some other IT component.

Some Security Management aspects must be kept separate - such as a policy enforcement or audit - but the whole IT infrastructure benefits from ingration between most security functions and the other IT/network operations

Further, if you have a security operations center that is separate from a network operations center, this can lead to security decisions that are made without regard to business impact or IT decisions that violate security policies and leave an organization open to attack. For example, what if an intrusion detection tool sounds the alarm and a security decision is made to shut down a vulnerable server? Sounds like a sensible thing to do. But what if that server happened to be in the middle of a business critical transaction completing a huge order? Wouldn't it be smarter to apply alternate mitigating controls and delay shutting down that vulnerable server until the huge order is processed? You would want to know if the risk is acceptable to the business. This example attempts to illustrate the interdependencies between security and the rest of IT management. Proactive security management can have a large impact on the business, it relies on an integrated IT management approach and the ability to see and respond with the big picture.

1.4. Maintaining Acceptable Security and Risk Levels

Perfect security is believed to be unattainable, and experts recommend spending only as much money as necessary to obtain the appropriate level of protection. The common question is “How much security is enough?” The answer depends on the result of a risk calculation that factors in the value of the protected assets, the known and reasonably anticipated threats against those assets, and associated vulnerabilities. Security management, in this sense, becomes a tool for managing risk.

Maintaining an acceptable level of risk is the highest-level business goal for proactive security management. The acceptable level of risk, however, varies for industries, organizations, and companies, and a functional proactive security management solution provides the correct levels of confidentiality, integrity, and availability to meet the individual organization's acceptable level of risk.

2. Purpose

The purpose of proactive security management is to protect business assets, enable business processes, and drive security costs down. To serve this purpose in a cost-effective and efficient manner, proactive security management is driven by several requirements:

- Protecting against evolving threats
- Enabling evolving and flexible trust models
- Combating increasing process complexity and related expense and manageability challenges
- Remaining compliant with internal security policies, applicable laws, and changing regulations, including the regulatory complexity involved in international transactions and business relationships

2.1. Protecting Against Increasing Threats

The threat environment is increasingly complex and rapidly evolving. Security incident reports are rising in frequency, viruses and other attacks are spreading at faster rates, the complexity of attacks is ever more sophisticated, and relatively sophisticated tools for unsophisticated attackers (so-called "script kiddies") are widely available. This environment leads to a number of security management challenges.

As the number of incidents increases and the nature of the threats constantly changes, distinct protection technologies to prevent new attacks are deployed - for example, firewalls, anti-virus tools, and Intrusion Detection or Prevention Systems (IDS/IPS). In many ways, this can be viewed as an ongoing "arms race" where the attackers come up with a new attack and the security industry comes up with a new defense.

And this "arms race" against attackers is not limited to certain classes of attacks such as worms or viruses. Attackers are studying the entire software stack from hardware to human, including physical hardware, operating system, applications and the behavior of human users.

Not only are there more and different kinds of threats and vulnerabilities, but the speed of attack has decreased to milliseconds from what used to take hours and days to spread. Old security management processes that require humans to respond with pagers, discussions and human-speed decisions, now must be enhanced with automatic systems that react in milliseconds to mitigate or slow down the fastest attacks.

Proactive security management is here to deal with this increasing threat environment with the latest set of technologies, methods and processes.

2.2. Enabling Changing Trust Models

The opening of business and organizational boundaries has changed old security models. With any combination of partnerships, mergers, dynamic supply chains, online customer services, federations, and changing user populations, it is very difficult to draw a line showing where an organization's intranet stops and the Internet begins. The old concepts of inside (people inside the organization employees or contractors) and outside (everyone else) no longer hold. The reports of incidents involving insiders show that this old, single-wall model of security was grossly ill-conceived and is not adequate.

Proactive security management must now protect a larger set of users that change over time, including a changing set of privileges based on roles and a set of resources that can expand and contract. This protection must match the speed of the changes. For example, when an employee joins or leaves an organization, access to resources must be enabled or disabled in a reasonably short time.

Globalization of organizations also has created a challenging set of changes for proactive security management to control: disparate organizations - whether two companies partnering or coalition military forces - come together and need to share specific, sensitive IT resources based on particular trust relationships, for a limited period of time. You can imagine how gaps in such relationships can prevent joint operations or even worse, compromise sensitive IT resources.

With the evolution of trust models and boundaries, one newer field of security management has been maturing more recently - managing access to the network. This field is a critical tool in protecting the trust boundaries at the network. This is discussed in more detail in the Network Security section of the Trusted Infrastructure chapter (Chapter 4) of this handbook.

2.3. Managing Increased Process Complexity

Each new protection technology or component introduces additional complexity in the organization. As new security technologies are purchased, they must be integrated and managed with the other existing security technologies/processes deployed in the environment as well as integrated within the large IT infrastructure.

For example, a corporate perimeter might use firewalls, routers, and gateways-each with a complex set of rules to create and maintain. Behind that might be some bastion hosts, which are server-class machines that provide Internet services and serve as a buffer or demilitarized zone between the open internet and an organization's private intranets.

Other components of a company's IT infrastructure might include an IDS, an IPS, honeypots and honeynets (servers or network segments acting as traps for attackers), an anti-virus program, and a security patch management system. From this quick example, the complexity of managing security technologies becomes apparent.

Correlating alarms and alerts, consolidating control, centralizing the reporting and management of the entire security operation, and developing in-house expertise for each of these components are significant, costly challenges.

Proactive security management solutions and services can simplify or largely offload the burden of this complexity (if outsourced or out-tasked), by providing functions such as:

- Consolidation of security information from multiple sources for a central view of security state of networks, systems and other IT resources
- Simplification of the consolidated information to facilitate higher level interpretation and decision making
- Decreasing costs of staffing specific tool expertise and freeing IT Security staff from more mundane tasks (like maintaining firewall rule sets).
- Enable higher-level controls on the specific, lower-level security tools that would allow policy based decisions to be automatically propagated and decomposed into instructions and configuration commands for the lower-level tools. For example, policy decisions about which IT resources can be utilized by which employees would automatically trigger configuration changes for routers, firewalls and anomaly detection systems to allow the correct access - no human work involved, and all done in a systematic, accountable process.

With such proactive security management capabilities, the security infrastructure gains the increased protection of new tools and methods, in an agile and response fashion, while maintaining compliance with security policy and a consistent level of risk.

2.4. Complying with Changing Regulations

Problems with privacy violations and lack of security generate press and public attention and cause far-reaching changes in the way we interact with governments, businesses, and organizations. Legislative bodies, standards organizations, and industry-specific groups have created laws, standards, and certifications to guide or mandate how organizations create, store, use, and communicate information.

In the U.S., for example, the Sarbanes-Oxley Act requires public companies to show that they preserve the integrity of corporate financial information and take steps to protect that information from unauthorized access. Another U.S. example is the Health Insurance Portability and Accountability Act (HIPAA), which requires enterprises to take meaningful steps to preserve the confidentiality of customer/patient information. Controls such as these drive the functionality of the security infrastructure and require proof of compliance by methods of auditing and event logs. In the EU and Japan, there is a lot of progressive work on privacy legislation and the security of personal information.

With an established set of minimum functionality, Security Management, like IT Management, must now mature beyond security and deliver business objectives.

Security management is the control point for the collection, transportation and storage of sensitive IT information and it is the focal point when things go awry or when the auditors come around.

Proactive security management addresses this responsibility with its controls and monitors which watch for variances, test for compliance, and provide audit and logging controls to meet laws and regulatory requirements. Further, security management will become truly proactive with constant compliance capabilities to eliminate the time-consuming preparations for audits and flag new configurations that violate security policies.

2.5. Purpose of Proactive Security Management Depends on More Than Technology

As an aside, it is important to note that security technology alone is not the only piece to achieve the purpose of proactive security management. In order to balance security technology, people and processes must not be overlooked. Continuously reinforced awareness programs and ongoing end-user training are essential. The more end users can learn about the actions they can and must take to mitigate threats, the more secure the enterprise will become. Also, the more that the enterprise can capture and learn from responses to threats and attacks, the more secure the whole enterprise will become. Unknowing actions can undermine the best-managed security infrastructure.

2.6. IT Management Trends and Security Management

With so much technology available, some argue that lots of technology is used for technology's sake and has lost sight of the goal: achieving business results. The shift from "Information Technology" to "Business Technology" changes thinking from information technology as a separate department to a model where technology powers the business. With it, the role of the CIO changes to be measured on overall business outcomes, rather than delivering only on technology service level agreements.

For security and privacy management this Business Technology thinking means that security management must focus on enabling business and organizational goals. For example, instead of a focus on the number of attacks repelled each month or the number of viruses detected, the CISO must now create metrics that demonstrate business outcomes. Table 2-1 has some examples of taking a business technology approach in terms of Risk Management and Compliance (rather than security for security's sake):

Table 2-1
A shift in IT thinking will change IT security program goals away from pure security or technology goals to those producing business outcomes.

Business Outcomes are Security Goals
• Confidentiality, integrity and availability of IT assets
• Ensure security controls investment appropriately matches asset values
• Optimize utilization of security controls and safeguards
• Reduce cost of achieving regulatory compliance
• Leverage security as a business enabler and market differentiator
• Achieve best-in-class security program to improve organizational reputation

3. HP Proactive Security Management Framework

There are many pieces to the security management program, including the process and people's roles as well as the security and IT technologies which provide specific security management functions. In figure 2-1, we present a top level perspective to give the reader a framework showing how individual pieces fit to form the whole program. In general the diagram depicts two rings - an *inner* and *outer* ring. Categories on the outer ring generally have much more interactions with the world outside an organization; the inner ring's categories generally stay more internally focused to an organization. There are color coded and underlined titles to designate which pieces are generally considered reactive security management (yellow) and those that are regarded as proactive (white underlined). Both proactive and reactive components are required for the most complete security management program. This circular framework diagram is chosen specifically to show that there is no hierarchical or chronological relationship between categories - each serves a specific function and covers its part of the problem space. There are definitely overlapping tools and functions that serve in multiple categories of this diagram; there are often tight links, interactions or dependencies between categories as well.

Figure 2-1
High-level view of the categories or pieces that make up the proactive security management framework

3.1. Compliance, Security Monitoring and Reporting

The evolution of the people, processes, tools and technology to protect your infrastructure is analogous to creating a quilt. Over time, independent pieces of security technology and tools have been invented to cover specific aspects of securing an infrastructure or to address a new family of threats or vulnerabilities. Many security infrastructures add the new technology/tool to their existing architectures as soon as practical, integrating it with existing systems as possible-creating an ever larger "quilt" that cover's IT infrastructure from a growing list of threats and vulnerabilities. Figure 2-1 shows the main categories that come together to form the latest set. With this growing set of tools in place, organizations are now receiving security information from each tool, and easily become overwhelmed: What do I do if one tool reports I'm safe and another tool reports an alarm? Do I have to staff a team for each tool? How do I make sense out of what my firewalls are reporting compared with what my intrusion detection system is reporting? Why does my anti-virus tool report all is fine when my networks are clearly under attack?

So emerged the need to coordinate and correlate information from different security tools. First came the correlation tool products that would receive alerts and information from different security tools and correlate the information to present in a single console. The correlation tools then evolved to gather and present tracking and workflow information so you could not only identify an event, but also track the mitigation efforts and get status information as the threat and mitigation work progressed. This current set of tools is labeled Security Information/Incident Management (SIM) or Security Event Management (SEM).

The current trends for SIM/SEM solutions are:

- Expand the scope of information collected to include integration with existing IT logs that have been collected for configuration and performance management. IT logs contain security-pertinent information about patch levels, system configurations, accounts and audit records.
- Reporting the status of compliance to specific customers set off by policies, laws and regulations. As organizations work to respond to current and future laws and regulations that require security and privacy mechanisms, it is very efficient and effective to have a security management architecture that can constantly check its own compliance and report on that status for audits when required.

3.2. Vulnerability Management

A security vulnerability can be defined as a software bug or flaw that can be exploited deliberately for unintended results. The field of Vulnerability Management has emerged in security management to focus on identifying such vulnerabilities or holes in existing IT infrastructures and managing their remediation or mitigation. There are tools that perform a series of tests that try to detect each known vulnerability and create a summary report stating how many known vulnerabilities are likely to exist. These tools are referred to as scanners, penetration testing tools or vulnerability assessment tools. They depend on advance knowledge of what a vulnerability is and how to deduce if that vulnerability exists. They use a database of known vulnerability descriptions and produce reports which state results in quantitative terms such as "System was scanned for 4,000 known vulnerabilities. Results: 106 vulnerabilities discovered with a priority breakdown of 27 critical, 17 high, 42 medium, and 20 low priority".

The current trends for Vulnerability Management solutions include:

- Reducing the number of false positives and false negatives
- Integration with the other IT functions of configuration management and system maintenance by including status reporting and workflow functions
- Improving up-to-the-minute vulnerability information in vendor-provided, known-vulnerability databases, by utilizing other external intelligence sources such as vulnerability research companies or vulnerability sharing groups
- Expanding dynamic attack methodologies and/or fuzzing techniques which are used to discover previously unknown vulnerabilities
- Combining vulnerability status reporting with compliance reporting

3.3. Content Management

As more and more of an organization's critical information and intellectual property is created and stored in an electronic form, the security properties (confidentiality, integrity and availability), as well as the privacy properties, become increasingly important aspects to manage. Also there is a lot of electronic information that flows into organizations that must be managed carefully to again meet security, privacy and business conduct policies, laws and regulations.

Examples come from many different parts of a business, such as controlling who gets access to a company's intellectual property (IP); ensuring that a sensitive document is not exposed in violation of a security or privacy policy; protecting critical financial data from accidental leak or intentional theft; blocking illegal content from entering your IT infrastructure; controlling unsolicited e-mails (aka "spam" or "junk mail") from entering your e-mail system and adding unwanted costs for storage, bandwidth consumption, and wasted employee productivity.

The need for organizations to manage their electronic content that flows in and out of their company and IT infrastructures has created the newer field within Security Management known as "Content Management".

A majority of electronic content has primary intended value and secondary significance beyond its intended purpose: e-mails can become evidence, files can become Intellectual Property, Instant Messages can become a liability, a transaction can become sensitive private data.

There are products now available in categories such as:

• Rights Management and IP Protection: these tools maintain/enforce property, copyright or usage rights to the electronic "property", such as IP, specific data, or software. Such tools are being initially deployed to control access to music, movies, corporate data, individual users' data, software programs or industrial secrets.
• Content Filtering: this is a set of tools that sit on networks, data centers and/or user's compute devices and watch what's being sent and received.

Depending on an organization's security, business or privacy policies, as well as laws and regulations, some content is illegal or in violation of some policy and must be carefully controlled, such as pornography, or violent or racist materials. Some network traffic content might contain attacks such as virus, worms, or other undesirable problems that are being spread over network protocols or at a higher level in web applications, such as web browsers. An emerging set of products called "Content Filtering" solutions will search/filter the traffic by signatures, patterns, keywords, behavior or other means to look for the sensitive content traveling in violation of policy, law or regulation. These solutions allow you to block, delete and/or report such violations. Products in this field are known by other names such as "Spam Filtering" or "Content Protection". Products that filter web content going in and out of web browsers are referred to as "Web Content Inspection" or "Web Surfing Controls".

3.4. Identity Management

Identity Management is the field of security management that encompasses the processes, tools, social contracts and the lifecycles of digital identities. It is such an important aspect of security that it has arisen as a large field of security with many technologies, products, methodologies and applications. For HP's overall security strategy, we have called out Identity Management as a separate pillar; subsequently we have dedicated an entire chapter to it in this document. For more information, please refer to Chapter 3.

3.5. Host Management

In our Proactive Security Management Framework, "Host Management" refers to the field of installing, controlling and maintaining the configurations and controls that have security and privacy implications and are included with operating environments and applications on various computers and devices. Such controls and configurations include file permissions, access settings, security settings, security-related patches, a computer's services availability, network settings, operating system parameters and drivers.

As you might imagine there are many different aspects to managing the security of a host system and much progress and innovation has been done in this field. Specific security market segments have emerged with tools and solutions not only from the operating environment vendors but from independent product companies as well.

Categories of Host Management aspects include:

• "Secure out of box" or "Secure by default": These phrases refer to a paradigm shift by manufacturers to send their products out with security-risky features turned off by default instead of making users actively reconfigure a product to be safe.
• Hardening: To make a system "hardened" means to close all the known holes and turn off all unnecessary services and accesses. There are tools that check and set network, device and system settings to known safe settings - like closing open network ports that are unused or removing device drivers that will not be used - thereby reducing the amount of avenues for attack potential.
• Chain of trust: Building a chain of trust for a computer system refers to creating a way to start from a known trusted item and to add additional functions (or software) on top - each depending upon the previous trusted item. For example if you wanted to know that your payroll program has not been tampered with, you would need to know that the operating system running your payroll program has not been tampered with. To trust the operating system, you would need to know if the boot code was not tampered with, and that the hardware has not been tampered with either. This would be a chain of trust between the hardware/boot code and your payroll program. Creating and maintaining such a chain of trust is another part of managing host security. Trends in building chains of trust and creating a trusted starting point are being led by the work of the Trusted Computing Group (TCG), which has developed an industry standard for Trusted Platform Modules (TPMs).

• Patch and Configuration Management: If the highest risks of attack come from unpatched software and misconfigured systems (e.g. a system configured with no password on a root account); then you can imagine how important testing, applying and verifying security related patches and configurations is to the security of a host system. This importance prompted a growth in tools and solutions to do this patch and configuration management in a systematic and auditable fashion. As this aspect of Host Management has evolved, it now links into audit and workflow systems to insure that changes are made, auditing can be done efficiently and effectively, and compliance to policy can be collected/reported.

• An emerging field in Host Management is the concept of performing system health checks to determine if a host meets a minimum set of security policies before letting that host connect to IT resources, like an intranet. There have been several leading methodologies proposed, including Trusted Network Connect (TNC) from the Trusted Computing Group, Network Endpoint Assessment (NEA) from the Internet Engineering Task Force (IETF), Network Access Protection (NAP) from Microsoft and Network Access Control (NAC) from Cisco.

The above aspects of Host Management illustrate a management view of securing a host. Host Management is one part of the bigger topic of Trusted Infrastructure, which is covered in Chapter 5 in this handbook. Please refer to that chapter for much more detail and a complete discussion of these Host Management aspects in the context of building trusted infrastructures.

3.6. Intrusion Detection and Prevention

The last element of the outer ring of the Proactive Security Management framework diagram (Figure 2-1) is the field of Intrusion Detection and Prevention. This field is focused on defining what constitutes "good" behavior on a system or network and assuming that any behavior observed which falls outside the definition of "good" must be mitigated. The definitions of good and bad behavior come from the particular type of detection and prevention methodology being examined as there are several aspects to this field that have been developed into product categories in the security industry.

Here are a few examples:

• Network Intrusion Detection and Prevention Solutions (NIDS or IDP): Products in this category typically run on network elements (e.g. a network switch or router) or dedicated systems connected to managed networks. They examine the network traffic for patterns of known, undesired behavior, or for anomalous behavior. They can react with alarms, alerts, logging and/or predefined reactions, such as dropping known-bad packets.

• Host-based Intrusion Detection and Prevention (HIDS or IDP): These products are similar to the network IDP, but instead of examining network traffic, they focus on the behavior of the host computer where they run locally. They look for known bad behaviors, like accessing protected files or many login attempts, and also look for anomalous activities. These products can react with alarms, alerts, logging and/or predefined reactions, such as locking accounts or blocking access to a resource.

• Anti-virus products: These were the original technologies that looked for bad behavior. They use information (in databases, signature files or rule sets) of what "bad" is to identify attacks such as viruses and worms. They have also evolved with the IDP category of products to look for patterns of suspect behavior as well as to scan whole file systems for infected files or other telltale evidence of attacks. These products will send alarms and alerts, log events and can take predefined reactions, such as blocking traffic from a suspect source, deleting infected files or moving suspected bad files into a special quarantine area to prevent usage.

• Firewalls*:* These were the first and most popular intrusion prevention products to be deployed because they originally restricted access to access points. Firewalls sit between networks and systems, examining network traffic and isolating traffic that violates security policy. Evolving from that initial simple, but very effective function, firewalls have increased the level of traffic inspection to examine more than just the connection request.

These examples above give a perspective of managing intrusions in this chapter's context of Proactive Security Management. However, intrusion detection and prevention solutions are integral in the architecture and construction of an effective security infrastructure and therefore are presented in the Trusted Infrastructure chapter of this handbook. Please refer to Chapter 4 for more detail and discussion.

The following categories of Proactive Security Management (Figure 2-1) make up the *inner* ring of the diagram. This inner ring represents those fields of Proactive Security Management that are generally more internally focused in an organization or company. There definitely are external links to the people, processes and technologies that are deployed for these inner-ring categories, but the focus still is more internal than external.

3.7. Problem Management

The main purpose of problem management is the rectification of errors in the IT infrastructure; its goal is to proactively minimize the impact of security issues on business and to prevent recurrence. A problem is often identified on analysis of incidents, which have some commonality in symptoms. However, problems can also be identified by analysis of a single resolved/closed incident of high impact with a possibility of recurrence. In either case, a business case would then exist to justify the expenses that accompany root-cause analysis associated with problem management. This is reactive problem management.

Problems can be identified by analyzing the IT infrastructure/reports, by using knowledge databases, interaction with developers/vendors on known errors when new products are launched, as well as meetings with the user community.

Once problem is identified, efforts are made to arrive at the root cause. Successful analysis of this root cause identifies the "Known Error" condition in the IT infrastructure. From that point forward, some kind of corrective action is defined and executed through a controlled Change Management process.
How are Problem Management and Incident Management different? Problem Management focuses on arriving at permanent solutions to known errors in the infrastructure. The objective of Incident Management is to restore normal service operation, often through implementing workarounds such as a temporary fix or routing the service to the customer through another Infrastructure channel.

3.8. Investigations and IT Forensics

When events occur in violation of security and privacy policies, laws or regulations, part of managing a complete response is to discover what happened and why it happened. This the mission of the Investigations and IT Forensics category and is made up of combinations of people, policies, processes and tools that are used to log, gather and present data in an investigation of security or privacy policy violations to answer the questions about who did what, when, how and why.

As with most of Security Management, there are reactive and proactive components to the IT forensic tools available for investigations:

- Reactive IT forensic tools are used during or after an event. They are used to collect evidence of the event in question and can include everything from a forensic system used to image hard drives, or a camera to photograph physical evidence to forensic software products that perform imaging, password cracking, decryption and specific evidence preservation processes.

- Proactive IT forensic tools and methods capture and protect forensic data that exists before and during a security/privacy event. These are called proactive because they are active before events occur and therefore are gathering data before and during when an event of interest occurs. When an alarm is triggered, identifying an event to investigate, proactive forensic tools provide additional useful data. Such proactive IT forensic tools and methods include both dedicated security forensic products and general IT system functions such as system event logs.

3.9. Security Program Administration

Administering an organization's security program is a general category referring to the activities that are required to create, run, enforce and maintain the security program's policies and governance models. On the governance side for example, authorization policies must be created, documented, communicated, enforced and audited. Such authorization policies might stipulate who is allowed to access which IT resources, who is allowed to grant access, and what process must be followed to grant and revoke access rights. Security policies for an organization could define the acceptable use of IT assets or the required security configurations for a system connecting to a managed network.

An adage that is repeated in security is that "security is people, process and technology…in that order!" The key message is that technology alone will not create a complete security architecture, without a well-defined process and trained and willing people using it. Security Program Administration includes the required education and awareness components. Even if the best security technology is in place, it can be instantly nullified by a person writing down a password on a piece of paper taped to a monitor. Also, a wonderful security process to protect sensitive company data will never be utilized effectively if less than half of the employees know about it.

3.10. Incident Management

As soon as an unintended security- or privacy-related event is detected, it can be identified as a security incident and Incident Management resources and processes must be activated. The objective of Incident Management is to restore normal service operation, often by implementing workarounds such as a temporary fix or by routing the service to the customer through another infrastructure channel. Often teams are staffed to execute the Incident Management procedures and are referred to as the Security Incident Response Team or Computer IT Security Response Team. These teams have an important responsibility to restore normal operations while at the same time preserving forensic evidence and operating within the boundaries of their organization's policies as well as other laws or regulations. Therefore they must have a predefined set of procedures that take into account business priorities, policies, laws, options for defensive techniques, and organizational escalation processes.

3.11. Risk Management

Security Management is one tool to manage an organization's risk; it is typically used in conjunction with other tools, such as insurance or outsourced services, to achieve an acceptable level of risk.

As was discussed in the introduction to this chapter, security has matured to a point where it is ready to move from an independent IT security operation to an integrated IT component that serves to support desired business outcomes. The examples earlier showed how security objectives will be constructed to deliver business results such as cost effectiveness or maximizing utilization of security infrastructures while maintaining compliance to organizational policies for security and privacy. When security is managed in order to achieve business objectives, security evolves from simply protecting IT assets to become a tool for managing risk.

It is difficult to measure the return on investment for a security infrastructure, as its goal is not to produce a business value but to enable other business activity. The true value of security comes from its ability to minimize or mitigate the risk interruption to business services. Since it is easy to overspend on any IT infrastructure and security is no exception, how much security is enough and how much is too much? The answer comes from the recognition that a security infrastructure allows an organization to manage the level of acceptable risk. Installing all the available security tools and purchasing insurance for unforeseen events would be very expensive and fiscally irresponsible. At the other extreme would be to do nothing - which would expose an organization and its officers to risk of damage, loss, litigation or prosecution. Further, there is no single security infrastructure architecture that is correct or best for all organizations.

Therefore, the optimal way to manage a security infrastructure at a higher level is to create a security Infrastructure that provides a constant acceptable level of risk. Measuring, tracking, auditing and reporting that risk is part of the Risk Management piece of the Proactive Security Management framework.

Risk Management is a new field for the security industry and many models have been proposed but none have achieved broad acceptance nor usage as of this writing. Risk Management will be a higher level of Proactive Security Management that will take into consideration such factors as business impacts, security risks, vulnerability states, and an organizations true appetite for risk.

3.12. IT Administration Integration

The trend for future security management is to integrate with the other existing IT management disciplines and give IT infrastructures the desired attributes - functionality, performance, availability, integrity, confidentiality, trustworthiness and reliability - all managed at the same level of corporate management. For example, consider the following three IT events: a critical server running out of disk space, the launch of a new sales portal resulting in a flood of online orders, and a security breach resulting in the theft of confidential data. What if all three of those events were happening on the same server machine at the same time?

If security management is completely separate from other IT functions, you might end up with the following responses to the above scenario:

- Security management immediately isolates the compromised system to contain the breach from spreading further and freeze all forensic evidence to establish a chain of evidence for the investigation…but this would immediately halt the revenue stream from the new sales portal.
- IT Administrators would add disk capacity and reconfigure the server to handle additional capacity…which might destroy forensic evidence and delay the identification and mitigation of the attack, and would also risk continued theft.
- Network management sees the spike in order traffic and immediately opens more network ports to keep the orders coming in…instantly exceeding disk space on the server.

This example shows that a perspective above security management, IT administration and network management is needed to be able to assess the current state of risk and urgency and to make a responsible, prioritized response to such a scenario. This is the case for bringing security management into the same place as the other IT management functions, consoles and operation centers. There is some work to do to literally integrate the plethora of tools and products and it is happening slowly. The guidance and leadership on how to integrate the lower-level, technical tools will come at the architectural levels.

There are several models that are emerging as pragmatic and widely accepted. One example is Information Technology Service Management (ITSM) with Information Technology Infrastructure Libraries (ITIL).

• ITSM*:* This approach combines proven methods such as process management and known industry best practices to enable an organization to deliver quality IT services that satisfy business needs and achieve performance targets specified with service level agreements.
• ITIL*:* This integrated set of best-practice recommendations is used to aid the implementation of a framework for ITSM. This framework defines how Service Management is applied within any type of business or organization that has a reliance on IT infrastructure. ITIL covers areas such as Incident Management, Problem Management, Change Management, Release Management and the Service Desk.

By taking such a top-down approach in architecting an entire IT infrastructure including security management, an organization will have higher level visibility across the whole IT infrastructure to make the best decisions to support the desired IT attributes of functionality, performance, availability, integrity, confidentiality, trustworthiness and reliability.

4. HP Proactive Security Management Offerings

Security goals, risk profiles and IT infrastructure maturity levels are unique to each organization. HP's proactive security management products and solutions have been created to enable a modular approach to thoroughly customize proactive security management solution components to meet an organization's specific security needs and budget. The primary elements of HP's proactive security management offerings are the HP Proactive Security Management services delivered by HP Services and proactive security management products provided by HP and HP partners.

4.1. HP Proactive Security Management Services

HP Services has a comprehensive portfolio of Security Services to help commercial companies and organizations establish and deploy a Proactive Security Management program or design a complete Security Operations Center. These are the Proactive Security Management Core Services based on the HP Proactive Security Management framework presented in the previous paragraphs.

A Proactive Security Management program takes into account strategy, people, processes, tools, and technology in a holistic and coordinated manner. The program helps companies proactively manage information security threats, vulnerabilities, and incidents in order to reduce their impact on the organization.

In addition to the HP Proactive Security Core Services, HP Services offers:

• Quick Security Assessment and Health Check Services which are linked to the Vulnerability Management component of the framework. For a comprehensive Security and Risk Management assessment, please refer to the Governance and Compliance chapter of this publication: Chapter 1.
• Managed Security Services for the HP Services IT outsourced customers from the HP Security Operation Centers, or outsourced security monitoring with the Enterprise Security Partnership service.

4.1.1. Security Assessment Services

4.1.1.1. Custom Security Assessment

The objective is to help customers establish and deploy a proactive security management program to effectively manage threats and vulnerabilities while minimizing the business impact of IT security incidents. Preventing security incidents from occurring requires proactive steps. In addition, for those incidents that do occur, a proactive security management program should:

• Restore normal service quickly and efficiently, with as little impact to the organization as possible
• Ensure that all security incidents are identified and processed in a timely and consistent manner
• Prioritize and provide direct support services where they are needed most
• Provide accurate information about the security incidents that occur to better plan and optimize existing security systems
• Identify, address, and correct or minimize any damage to systems or data
• Evaluate the effectiveness of response(s) and feed a knowledge management system (if available) to learn what worked well and what did not

Benefits

With a comprehensive proactive security management program in place, organizations can:

- Minimize downtime, exposure, and loss of critical information caused by security attacks, thereby minimizing damage to business, company brand, customer loyalty, intellectual property, and employee productivity.
- Make security incident and crisis management decisions based on real-time assessments of threats and vulnerabilities, with an associated audit trail and action record to validate proper response and derive strategies and tactics for improvement.
- Prevent or minimize the spread of security attacks within the enterprise and stop the propagation of worms, viruses, and other pathogens.
- Control internal information for compliance with regulations (for example, Sarbanes-Oxley and the Basel II Accord) and prevent liabilities under the regulatory mandates.
- Focus on business rather than security incident recovery.
- Control security investments by focusing on the business impact of threats and vulnerabilities, thus informing relevant procurement decisions and ensuring maximum benefit.

Barriers

Minimizing downtime from security threats, vulnerabilities, and incidents requires a comprehensive response plan. Yet few enterprises have such a program in place. HP has found several real and perceived barriers. The most important barriers are:

- An overtasked security staff that is busy dealing with the mundane. The staff may also lack basic security automation tools such as patch management, group policy, configuration control, intrusion detection, and proper training or support resources.
- Missing or poor coordination across business units, enterprise application owners, data centers, and help desks. Fractured budget and investment policies discourage global, enterprise-wide investment in favor of local quick fixes. These practices result in an overlapping patchwork of solutions that cannot be managed effectively and may in fact work at cross purposes.
- The lack of a security governance model for prompt incident response decisions, making an analysis of the return on investment almost impossible to formulate and justify.

4.1.1.2. Core Services Associated with the Proactive Security Management Framework

At HP, we understand that organizations have an existing IT infrastructure and some processes in place to prevent and manage security incidents. Some of these processes are well suited to their purpose in the context of an integrated view of the system, and they can be carried through as part of a strengthened and proactive security management program. Some processes need to evolve or change, others should be eliminated, and new processes are inevitably needed. Deploying an effective proactive security management program is a transformation and an ongoing collaborative effort involving many facets of an organization.

The HP Proactive Security Program involves a deliberate process driven by HP to help organizations identify their needs and implement effective solutions. As Figure 2-2 illustrates, HP uses a best practices workshop approach to identify stakeholder needs, the current state, and the desired future state. An architecture and planning process, in close collaboration between HP and key stakeholders, results in the functional implementation of a proactive security management program.

Figure 2- 2
Overview of the HP proactive security management services

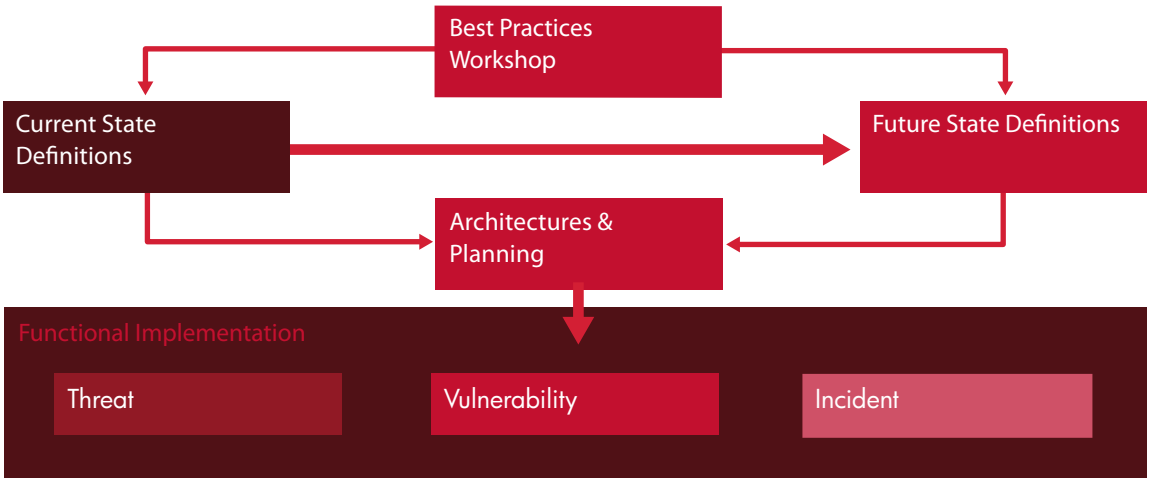

1. The *Best Practices Workshop Service* helps the customer understand the benefits of adopting a proactive security management solution approach.
2. The *Current State Definition Service* defines the current state of the customer's security infrastructure, process and organization in order to establish a baseline on which to build a proactive security management solution.
3. The *Future State Definition Service* defines the customers desired end state for proactive security management and a detailed roadmap to get there.
4. The *Planning and Architecture Service* involves the production of a solution architecture for a Proactive Security Management framework to encompass the specific future state defined for the customer together with implementation plans to support its deployment.
5. The *Functional Implementation Service* involves the design, implementation and integration of the individual proactive security management solution components required by the customer.

4.1.1.3. Methodology

HP has developed a highly structured process to design and deploy complex security solutions consistently and effectively - regardless of variations in organizations and IT architectures. The first phase of HP's methodology is assessment and planning. This is followed by the design and implementation phase.

Assess and Plan

The first phase of the process is to gather information, assess it, and define a plan. The steps include:

- Identify stakeholders and work with them to define the specific business value drivers for the organization.
- Perform a gap analysis between the current state of the IT infrastructure and the desired and appropriate future state.
- Perform a business impact assessment to prioritize the risk that each gap represents.
- Define a gap closure plan that identifies and prioritizes transformation projects.

Design and Implement

The second phase of HP's structured process is to design and implement the proactive security management program and associated tools. The steps HP undertakes include:

- Organize by establishing a project management structure, preparing project work plans, and setting objectives, milestones, and metrics by which to measure progress and success.
- Set expectations and timelines for program metrics reports.
- Design and document the program itself, and develop a handbook outlining roles and responsibilities, processes, and tools for implementation.
- Report vulnerability correlation and incidents and implement program area tools.
- Create and deliver training to build awareness.
- Ensure that systems are in place to effectively capture lessons learned and facilitate a feedback mechanism. This helps organizations learn from experience and makes this information available to staff in a usable form. It also feeds the training program, as appropriate.
- Develop and implement a maintenance plan.

4.1.1.4. Summary

HP's methodology and structured process for proactive security management produces reliable and effective results. By seeking to proactively identify security threats, eradicate vulnerabilities, and rapidly respond to attacks when they do occur, the HP Proactive Security Management Program protects information assets, prevents application downtime, facilitates maximum network/system/application availability, and helps to greatly reduce annual expenses caused by viruses, worms, and other costly security incidents.

The complexity of a best-in-class proactive security management program is directly related to the size and complexity of a company's IT infrastructure, its geographic reach, and its business needs. Many of the functions of an effective proactive security management program are combined and owned by a few individuals within the organization. However, regardless of how many people are involved, the functions and processes need to be well defined and followed.

HP's Proactive Security Management Program has been very successful. As a testimonial, the same team that helped develop and implement HP's internal program also developed our worldwide consulting services offering. It is important to emphasize that the entire cost of developing and implementing the comprehensive HP Proactive Security Management Program can be justified by comparing it to the cost of damage from a single security incident.

4.1.2. Quick Security Assessment

The cost of preventing a security breach is always far lower than the cost of recovering from one. That's why HP provides expert services to evaluate your overall security strategy, identify the strengths and weaknesses of your current security posture, gauge the risks to your mission-critical IT infrastructure and business data, and show you how to address potentially damaging security vulnerabilities.

4.1.2.1. Custom Security Assessment

Receive an in-depth analysis of information security risks within your business-critical technology infrastructure. Tailored to your specific needs and environment, HP's Custom Security Assessment service takes a holistic approach to security management across multiple IT components. Coverage can include the status of your policies and procedures for security management according to the BS 7799 (ISO 17799) standard; the security posture of servers, storage, operating systems, applications, and databases; and the configuration and management of the physical environment.

Security deficiencies are uncovered through interviews with key members of your technical staff, audits of compliance with your security policies, configuration audits, and an onsite review of your physical security safeguards.

HP Services security specialists work with you to match the assessment's scope and level to your technical and business requirements. Deliverables are spelled out in a detailed Statement of Work. A final security report highlights threats and vulnerabilities, and offers recommendations for improvements.

4.1.2.2. Security Quick Assessment

HP's Security Quick Assessment Service gives you a convenient, cost-effective way to gain an awareness of potential vulnerabilities in your IT environment and get expert recommendations for remedying them. During a one-day workshop, HP security consultants direct key members of your staff through a facilitated security self-assessment. The consultants then analyze your responses and report back to you on any weaknesses in your security management systems, as well as suggested avenues for improvement.

Assessment criteria are based on industry best practices such as ISO 17799, plus HP's wide experience in security solutions design and support.

4.1.2.3. Security Vulnerability Assessment for Small to Medium Businesses

Small and medium-sized businesses (SMB) face the same security threats as enterprise companies. Security Vulnerability Assessment for SMB can help by proving accurate, actionable information to start building an effective security plan.

Security isn't just for large enterprises any more. In fact, a recent study by Forrester Research found that "75% of small and medium businesses (SMB) expect to make new security investments in the next twelve months."

There are two primary reasons for the growing interest in, and need for, security among SMBs. The first is the rapidly-growing number of viruses, Trojan horses, hacker attacks, and other threats that today target large and small companies and individual and corporate users alike. The second is the fact that businesses of all sizes now rely extensively on their IT environments to meet their business goals.

Unfortunately, designing and implementing a security plan that provides the best possible response to threats, while also ensuring the most effective and focused use of budgets and resources, is often a major challenge for SMBs. That's exactly where the Security Vulnerability Assessment from HP can help.

This service provides access to a trusted partner with industry certified credentials who can lead SMBs through the complex security field. In the strictest confidence, HP security consultants prepare the accurate, actionable information these companies need to identify specific network security vulnerabilities at the most affordable price in the industry. The service also includes expert assistance to help SMBs identify the remedial solutions that will deliver optimum results in addressing their specific vulnerabilities.

The HP Security Vulnerability Assessment for SMB service allows SMBs to obtain actionable information to start building a security plan at the most affordable price in the industry. The service has the following features and benefits:

• Comprehensive: Provides an impressive level of analysis, including penetration testing of perimeter systems, done by experienced HP Security Consultants.
• Proactive: Identifies your business exposure to today's IT security risks by locating vulnerabilities and weaknesses in your networking infrastructure before it impacts your business.
• Informative: Helps you understand your current IT security measures and how they compare to industry benchmark standards.
• Performance-proven: Leverages HP's 25 years of experience and expertise in creating and delivering security solutions to customers worldwide.
• Affordable: Designed and priced specifically for SMB realities and requirements.
• Fast and easy to purchase. All essential elements included in a single package.
• Flexible: Two service levels available - Basic and Enhanced. You select the one that's the best match for your environment and its requirements. The basic and enhanced service levels both provide security architecture and policy reviews, penetration testing of perimeter systems, wireless security reviews, discovery and recommendations reporting, and best practices sharing. The enhanced service level additionally defines a security patch strategy.

4.1.3. HP Security Enhancement Services

The HP Security Enhancement Services offer small and mid-sized companies a much-needed new approach to addressing the growing number of security threats that are now plaguing companies of all types and sizes. This approach lets customers purchase and use affordable HP Care Pack Services units to access the specific services they need to address their environment's vulnerabilities and their business' security requirements. HP Security Enhancement Services for SMB can be used:

• To quickly follow up on the findings and recommendations of the HP Security Vulnerability Assessment for SMB service. This assessment identifies exactly where a company is vulnerable and what it can do to best address those vulnerabilities.
• At any time to address any security-related concern or issue that requires specialized expertise.

HP Security Enhancement unit of this service is designed specifically to help small and mid-sized companies access the capabilities and solutions they need to minimize threats, while also building a more secure overall environment. It follows the recent introduction of the HP Smart Desktop Management Service, another unique service that offers a complete approach to security for multi-vendor networked PCs in one integrated, affordable off-the-shelf solution that any SMB can easily leverage.

4.1.4. Security Health Checks

HP Security Health Check Services provide quick yet comprehensive exposure and risk assessments, zeroing in on key components of your business-critical infrastructure.

An HP Services security professional consults with your IT team to identify the configurations, systems, or databases to be checked, then scans and analyzes them to uncover security weaknesses. Next, we prepare a detailed report outlining the results of the analysis and offering recommendations on how to address high-risk security vulnerabilities. Finally, we review the report with you and discuss a follow-up action plan.

You also have the option to deploy the scanning software for additional monitoring functions and more frequent scans by HP or your staff.

4.1.4.1. Intranet Security Health Check

This service provides a network-based vulnerability assessment of business-critical systems connected to your company's intranet - including key servers, network switches, and routers. It uses comprehensive, automated network security vulnerability detection and analysis to probe target systems and identify security holes. HP Services professionals help you understand your risks and identify the steps required to harden your infrastructure.

4.1.4.2. System Security Health Check

Focusing on the operating system level of your critical servers, this assessment uses a host-based approach to detect platform security weaknesses that are not visible to network scanning. Your system-specific security risks are identified, analyzed, and prioritized, and you receive expert recommendations for implementing appropriate corrective actions.

4.1.4.3. Database Security Health Check

Obtain vital information for improving data integrity, availability, access control, and security management. Databases used by your critical business applications are scanned for security vulnerabilities without affecting your production environment. A summary report outlines recommendations in areas such as authentication, authorization, and system integrity.

4.1.5. HP Managed Security Services

4.1.5.1. HP Services Global Security Centers

HP Services Global Security Centers provide defined and custom solutions in areas ranging from security awareness enhancement and security policy design to risk mitigation, security infrastructure development, security integration, and security training. They are located in locations such as Redmond, Washington; Hong Kong; and Grenoble, France. Figure 2-3 illustrates how the centers take information from many different sources to deliver their set of services.

Figure 2-3
HP Global Security Centers bring different sources of threat and vulnerability information together to deliver a set of services.

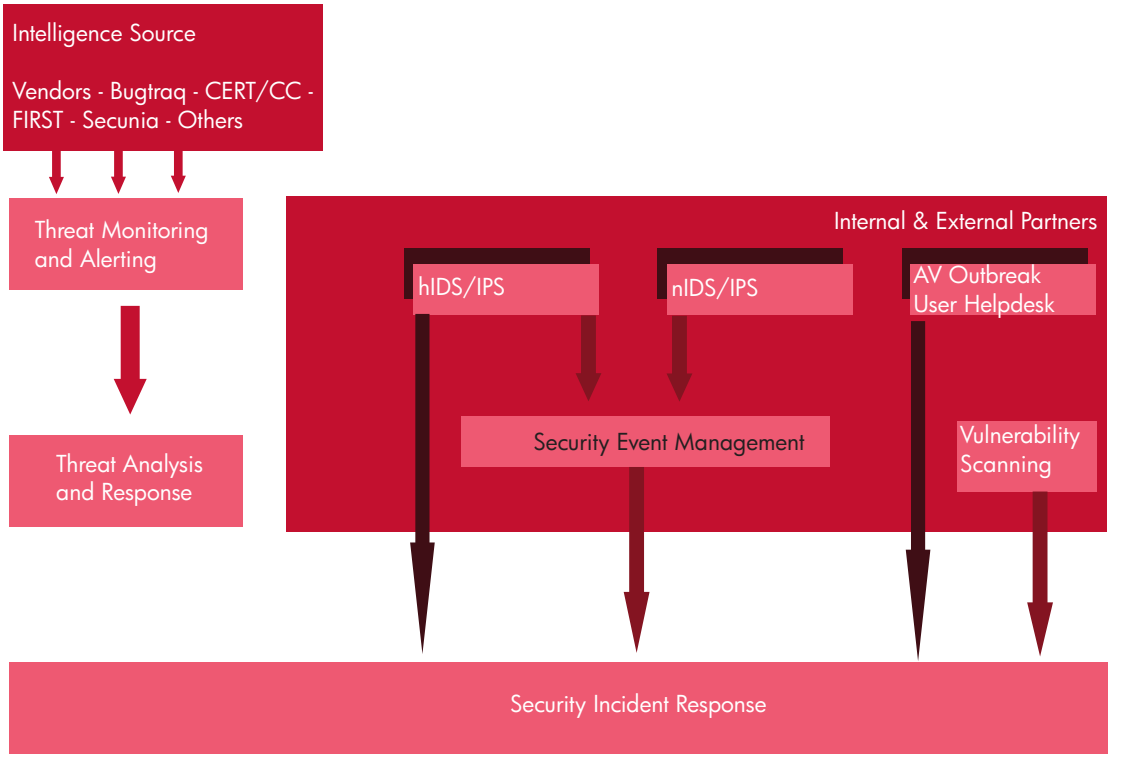

Services available from the Global Security Centers include:

• *Vulnerability assessments:* Numerous aspects of your infrastructure are examined to determine whether security gaps exist and how to correct them.

• *eSecurity probe:* Using the same tools and techniques commonly exploited by hackers, HP Services specialists can safely challenge the effectiveness of your perimeter security safeguards. You get an accurate picture of your preparedness to defend against a predefined set of common security threats.

• *Express services:* Fixed time and scope risk-mitigation services target specific aspects of your infrastructure, including firewalls, Web servers, Web applications, wireless networks, and telecom systems.

• *eSecurity scan service:* Vulnerability scanners combine with a proprietary HP methodology to provide a quick snapshot of your security vulnerability.

• *Incident handling service:* HP Services experts help you recover compromised systems in the shortest possible timeframe, identify the sources of attacks and take appropriate measures to prevent them from recurring, and track down attackers via forensics investigation.

• *Infrastructure security design service:* Deliver architecture design recommendations and network topology evaluation.

• *Security Training:* Customizable courses and complete programs facilitate your establishment and enforcement of ongoing security measures.

• *Pre-packaged security solutions:* Pre-integrated components, products, and services are leveraged to cut your time-to-results and keep your total costs down. Final integration and implementation are accomplished through custom services for your specific infrastructure.

Examples of solutions that can be provided by the Global Security Centers include:

• PKI solutions: These solutions typically encompass security policy, user registration authority, certificate issuing authority, certificate distribution system, and key archiving and renewal.

• Smart card solutions: Services include providing BIOS security, drive lock, and data encryption on client devices; personalizing smart cards with the HP smart card management system; and managing smart cards throughout their lifecycle.

• Enterprise Access Management Solutions*:* Make Web sites and collaborative relationships secure by providing centrally managed access control with distributed and delegated administration. Scalable to millions of users, these solutions support heterogeneous IT environments.

• Identity management and provisioning solutions: Provide a centralized means of managing identities, provisioning services, and implementing a consistent access control policy. Capabilities include user provisioning, self-service provisioning, self-service administration, and role-based policy management.

4.1.5.2. The Enterprise Security Partnership Service

HP Services offers the Enterprise Security Partnership service, which provides customers with a combination of world-leading IT services and security intelligence to address the ever-expanding security threat. This service is delivered via the joint expertise of HP and Symantec. It provides consulting, resources, and technology that minimize the disruption security attacks cause to businesses. The service helps organizations manage and defend their infrastructure, reduce vulnerability, and optimize the integrity and availability of the IT infrastructure as a key business asset.

The Enterprise Security Partnership service accomplishes this by providing:

• Real-time security monitoring and incident management
• Ongoing security management support
• Proactive security planning and improvement

The Enterprise Security Partnership service provides tailored services, including delivery of impact analysis reports for new and existing vulnerabilities and provision of proactive patches. Businesses also benefit from continuously improving security policies and procedures, regular reports on security threats and potential business risks, and effective countermeasure recommendations.

Symantec's extensive expertise within the security marketplace supplies industry-leading monitoring, intelligence, and analysis capabilities, protecting the infrastructure against attack. Combined with HP's security services and solutions and IT service management skills, organizations can improve security governance and implement and sustain a strong defensive posture to more effectively manage operational risk.

4.2. HP Proactive Security Management Products

To support the Proactive Security Management Program, HP offers a number of security monitoring and reporting products. They include solution components for incident response, security monitoring and event management, security reporting, patch management, and vulnerability remediation. HP (along with key partners) completes the picture with security event filtering, aggregation and correlation, policy compliance and vulnerability assessment, and various network security components.

Host Management: Patch and Configuration. HP Configuration Management solutions enable IT to respond to these demands through automated deployment and continuous management of software, including operating systems, applications, patches, content and configuration settings, on the widest breadth and largest volume of devices throughout the lifecycle for compliance by ensuring only authorized software is maintained on systems and policies are continually enforced.

Log Management. HP Operations Center provides a consistent system and fault management process and workflow. Intelligent agents can detect any security failure, and can monitor system and application log files for security problems, with the aid of partner solutions.

IT Administration and Integration. HP ServiceCenter helps IT organizations quickly deploy consistent, integrated work processes based on the ITIL framework. It is designed to help evolve IT service management organizations through a series of logical steps, from establishing basic controls to a higher level of maturation where automating service delivery can help maximize the business value of your IT organization.

Vulnerability Management. The HP Application Security Center, formerly SPI Dynamics, delivers a comprehensive and accurate suite of application security products and services that support the entire Web application lifecycle, from development and quality assurance to deployment, ongoing operations management and auditing. With over 5,000 unique Web application-specific vulnerabilities, threats, and security checks, SPI Dynamics' vulnerability database, SecureBase, is the most comprehensive and accurate knowledge base on the market. SPI Labs' hands-on experience in application penetration testing consulting, combined with extensive research to keep the database current, results in frequent updates.

Escalation and Crisis Management. The HP ServiceCenter Incident Management module automates the entire incident lifecycle, from the time a service disruption is reported through final service restoration. It can capture and log security-related incidents, including routing and escalation workflows based on criteria such as impact, urgency, or customer.

Security Event Management. HP Network Node Manager (NNM) provides the industry's leading SNMP-based environments, widely used by network-based security solutions to transmit security event information. HP Operations extends this with best-in-class system and application event management including log file monitoring and analysis. Event Correlation Solutions (ECS) provides a flexible mechanism for rule-based event correlation and processing the real-time event flows. Partners are also essential to the HP security event correlation and management strategy. Numerous integrations exist between individual security solutions and Operations Center for monitoring the security devices or applications and collecting the associated security events. This spans basic SNMP integration up through tested and certified Smart Plug-Ins (SPIs). For heterogeneous and comprehensive multi-vendor security event management, HP has partnerships with ArcSight, eSecurity, and Symantec. ArcSight and eSecurity provide certified SPI integrations for HP Operations. Symantec is in the process of completing a SPI integration. All of these partners provide security correlation capabilities that then forward correlated security events to Operations.

Compliance and Security Monitoring. HP Compliance Manager continuously monitors internal controls of key business processes, as well as their supporting applications and infrastructure, to measure effectiveness and mitigate risk. By aggregating and summarizing key metrics collected by HP Software tools and other control sources, HP Compliance Manager software shows the risk to high-level business processes, highlighting where there are control violations and emerging risk, enabling IT to quickly pinpoint and resolve the issue.

Configuration Management. HP Configuration Management (CM) Patch Manager software eliminates known software vulnerabilities quickly and reliably by automating the patch management process - including acquisition, impact analysis, pilot testing, discovery, assessment, deployment, maintenance and compliance assurance - ensuring that devices are always configured correctly. Using this policy-based software, IT managers can accelerate the correct configuration of their software infrastructure and optimize the security and stability of managed systems. The HP CM Patch Manager provides value for business continuity and security initiatives, server provisioning and repurposing, and OS and application migration. HP also offers a vulnerability and patch management solution for HP ProLiant servers as well as specific security-patching solutions for Microsoft Windows, HP-UX, and Linux.

4.2.1. Security Event Management

Security Event Management (SEM) is the ability to monitor and manage security across the entire IT infrastructure - from systems, applications, network elements, and security devices to all of the communications and transactions occurring within the infrastructure. SEM is a process embodied in the policies, network hardware, and specific SEM applications and services resident in the IT infrastructure. In total, it presents a complete view of the entire range of IT security elements.

4.2.1.1. Objectives

A comprehensive SEM solution actively records, views, analyzes, and manages all of the security events that occur within enterprise IT infrastructures. SEM includes the correlation of security data from multiple devices and systems across the enterprise to help facilitate security assessments and provide appropriate mitigation strategies and solutions.

An effective SEM solution aligns with the IT infrastructure so that security events can be judged in the context of the associated business risk. The capacity to determine the system's level of exploitability, an event's impact on business service(s), and the weight of assets at potential risk all contribute to determining the criticality of a potential security incident.

4.2.1.2. Environment

The major challenge facing the contemporary enterprise goes beyond its own borders and carefully controlled wide-area linkages. Enterprises are operating with multiple platform types and security products and services in an environment that exhibits an ever-widening array of connectivity requirements across partners, customers, and remote offices. Only a centralized view can identify incidents that require remediation and harden enterprise systems against future attacks. A centralized view aggregates all security events - no matter what, when, or where -and intelligently correlates the events with activity patterns.

Evolving government regulations and the regulatory challenges posed by multi-national operations are reinforcing the need for an effective SEM solution within an enterprise. Regulations now mandate organizations to implement security controls, and they hold organizations accountable, both legally and financially, for security incidents that compromise private information. These regulations drive architects and developers of IT security infrastructures to find a solution that constantly monitors networks for vulnerabilities. This has led to a number of toolsets, appliances, devices, and applications of increasing sophistication and scope, resulting in increasing complexity, integration, and management challenges.

4.2.1.3. Benefits

The correct SEM solution (properly architected, implemented, and administered) significantly eases the burden on overworked IT and security departments. By quickly identifying and responding to security threats and changing from a reactive mode to a proactive, systematic methodology, an SEM solution provides a productivity boost and reduces the direct costs of security implementations over time. The net effect is more efficient compliance with government regulations, protection of the corporate assets (and thus the bottom line), and smoother business operations.

By integrating a complete solution for SEM with the overall architecture of an incident management program an enterprise is able to:

- Insulate the higher-level incident management processes from the dynamic, ever-changing details of the security profile of a typical enterprise network-a network that is not only heterogeneous (containing a range of technologies, applications, and vendor-specific solutions) but also diverse in geography, time zone, use patterns, and languages down to the OS level.
- Conduct event filtering at the correct level while retaining sufficient audit and action records to validate any security oversights of critical components and information, if needed for legal or regulatory reasons.
- Manage the SEM system from a higher-level perspective and easily adapt the system to changes in local infrastructure and network conditions (granular adaptation) as well as to changes in the overall threat level and profile on the worldwide network (global adaptation).
- Demonstrate due diligence if a regulatory agency investigates the enterprise for compliance or if a legal action related to a security incident arises.

4.2.2. HP Software for Proactive Security Event Management

An HP Software-based SEM solution enables enterprises to detect and dynamically respond to changing circumstances. It also helps to securely manage evolving IT environments, minimizing operational impact due to security events or operational burdens arising from security solutions and methods impeding normal system usage. Intelligent partner integration, along with the best point solutions, provides an end-to-end global and local security management solution that proactively mitigates security incidents.

HP Software and security solutions from our partners can be integrated in to a broad Security Event Management (SEM) solution. Such an SEM solution is designed to centralize and manage all aspects of a security event.

Detection of a security event can come from multiple sources: for example, IDSs, firewalls, and system log file monitoring and analysis. Preventative notifications of potential security incidents, including unusual usage patterns and unapproved configuration changes, and early warning services can help mitigate incidents before they do damage. HP's unique Partner Integration Strategy achieves the highest level of integration in the industry through the development of Smart Plug-ins (SPIs) that integrate with HP Operations.

SPIs collect and intelligently analyze alerts. They correlate alerts as necessary, and forward them (if appropriate) to a higher level in the management hierarchy. SPIs also monitor the health, performance, and availability of the individual security applications and devices. SPIs are the preferred integration method for linking security devices and applications into HP Business Technology Optimization (BTO) Software.

HP's Event Correlation Solutions (ECS) can correlate individual event streams while also correlating events across security, system, application, and network sources. ECS provides the flexible mechanism for rule-based event correlation and processing the real-time event flows. For network-based security management, HP Network Node Manager provides the industry's leading Simple Network Management Protocol (SNMP)-based system, including log file monitoring and analysis.

HP's partner integration strategy achieves the following:

- Unified fault management covering all subsystems, collected and reported in a centralized fashion
- Unified reports covering specific incidents plus broader usage trends
- Configurable event filtering to shield operators from trivial matters, enabling them to focus on the most critical issues
- Event correlation to deduce cause and effect from seemingly dissimilar events
- Automated actions in response to a security problem, such as shutting down a process, paging an operator, generating an incident trouble ticket, or initiating a change management process (for example, to deploy patches)
- True service-level management, where specific security problems are immediately linked to those services they may affect, so actions can be taken in line with broader business objectives

4.2.2.1. What HP Provides: HP and Partner Solutions

HP Software provides a framework on which an enterprise can build a complete security management solution. A best-of-breed global management solution relies upon a collection of point solutions integrated into the HP architecture.

Creating the best of both worlds results in effective security management only if local solutions and global management are integrated to act as one. Such integration goes far beyond simply passing events from one application to another. Excellent integration takes full advantage of the local products' understanding of the managed object, plus the global solution's understanding of the complete infrastructure.

HP's approach relies on both internal expertise as well as that of our partners, and it provides enterprises with the ability to select the correct and most effective local solution for their situation.

4.2.2.2. HP's Partners and HP Software

Currently, HP Software offers SPI availability for the following leading security solution partners: Symantec, ArcSight, BindView, Check Point, Cisco Systems, e-Security, Netegrity, NFR Security, Nokia, Perfigo, Sun Java System Identity Server, Solsoft, Top Layer Networks, St. Bernard Software, ISS, and Tripwire. For an up-to-date list, see **www.hp.com/go/software**.

4.2.2.3. Security Event Management Summary

HP has built a complete solution for an enterprise SEM program that is exceptionally robust, sophisticated, flexible, and scalable. As a component in managing the Adaptive Enterprise, the HP solution is based on three critical elements:

- The scalable, secure, and proven HP Software platform
- A set of underlying data collection and analysis applications, running as SPIs to the HP Software platform
- Extensive and sophisticated processes and procedures that tie all components together with a robust and scalable platform and use additional HP Software components such as trouble ticket management to provide overall IT service management capabilities

The primary benefit of the HP Software approach is that an enterprise can take a holistic view of its entire IT infrastructure over the complete lifecycle of the infrastructure's individual components. As the networks, systems, and applications build and adapt to the changing requirements of the enterprise, so can the HP Software solution.

Security event management becomes integrated with the entire organization's approach to IT and network management. Components (such as firewalls, patch management solutions, and IDSs/IPSs) that once were islands of individual solutions and tactical approaches to local problems become part of a comprehensive solution.

4.2.3. Security Configuration and Patch Management

Security configuration and patch management incorporates security patches, correct configurations, and current versions of software. Eliminating vulnerabilities before incidents can occur is the greatest defense against attacks. By addressing a particular vulnerability, that threat is instantly and permanently removed. Security patch management fits with the compliance monitoring Solution component of proactive security management.

4.2.3.1. Objectives

Security patch management solutions should provide quick and reliable automation of the patch management process. Policy-based solutions help managers ensure that systems are current and that the security and stability of systems are optimized. Updating network nodes with the latest security patch is only part of the battle. Knowing what security patches to deploy, the effects they are expected to have on related systems and processes, and when to deploy them are key parts of any security patch management solution.

4.2.3.2. Environment

Although the patches are available, their deployment proves to be a challenge to enterprises. There are multiple reasons for this:

- Frequency of vulnerabilities and patches: The monthly rate of security vulnerability discovery and posting of patches has risen exponentially over the past few years.
- Quality versus speed during patch application: Testing and qualifying patches to ensure that they will not adversely impact the overall operation of the system has always been a concern. The speed in which security patches need to be deployed make this testing and qualification process particularly challenging.
- Ability to audit patch implementation: The ability to audit systems to ensure patch compliance has been an ongoing challenge, especially in diverse, multi-operating system environments with desktops, servers, and mobile devices.
- Number of mobile and remote users: Additional mobile and remote users bring challenges to an enterprise's patch capability. Mobile and remote users miss critical notification if they are not connected when an audit or scan takes place. Once notified, they may have limited bandwidth for accessing and installing patches. Finally, they may be exposed more quickly to attacks because they are located outside the protection of the enterprise.

4.2.3.3. HP's Patch Management Programs

HP Services combines the expertise of Certified Information Systems Security Professionals (CISSPs) and Microsoft Certified System Engineers (MCSEs) to recommend and deploy the right security patch management solution for an organization. Solutions cover a range of operating environments, including Microsoft Windows, Linux, and HP-UX. These solutions include HP Configuration Management, ProLiant Essentials Vulnerability and Patch Management Pack, Microsoft Security Patch Management tools, and HP-UX Patch Management tools.

4.2.3.3.1. HP Configuration and Patch Manager

HP Configuration Management solutions enable IT to respond to these demands through automated deployment and continuous management of software, including operating systems, applications, patches, content and configuration settings, on the widest breadth and largest volume of devices throughout the lifecycle for:

- IT efficiency to control management costs
- Agility to deploy services faster and without user disruption
- Security and compliance by continually enforcing policies, patch compliance and software integrity

HP Configuration Management software provides automation and control for every aspect of change execution with a suite of industry-leading tools as outlined in the following paragraphs:

Discover software and hardware inventory
HP Enterprise Discovery software is a part of the Configuration Management suite of products, for agent and agent-less discovery and inventory collection on hardware and software assets. Inventory, utilization and Windows Vista readiness reports can all be accessed directly through the Enterprise Manager console.

Streamline packaging and analyze configuration impact
HP Configuration Management Extensions for Windows Installer transforms any IT administrator into an expert in the advanced features of Microsoft Windows Installer. It has a unique, wizard-driven process that enforces best practices, streamlines the package building process, speeds troubleshooting and tailors the package to the needs of the environment. Impact analysis capabilities enable administrators to test for possible conflicts, to anticipate problems and help make new package rollouts run smoothly.

Speed PC provisioning and migration
HP Configuration Management OS Manager automatically provisions and maintains the right operating system for each device as prescribed by policies. It creates images, provisions them according to policies and manages the operating systems throughout the entire lifecycle. In a PC environment, it works together with HP Configuration Management Settings Migration Manager for personalized settings migration on each PC to ensure productivity during the migration process is maximized.

Deploy applications and content with ease
HP Configuration Management Application Manager provides the control and reliability required to execute timely application deployments based on business or IT needs. Everything can be handled with ease-emergency situations when patches or applications must be deployed immediately, scheduled deployments where the application must go live across the enterprise at the same time or small targeted deployments for a select group of end users. In a server environment, Application Manager utilizes Application Management Profiles (AMPs), which provide templates to ease the deployment of complex server applications.

Enable self-service software management
HP Configuration Management Self-Service Manager provides a self-service portal in a PC environment that users can access for downloading, repairing, updating and removing software. It presents a personalized software catalog via the electronic Definitive Media Library. This software library is generated dynamically according to the user's identity and role, the machine's configuration and the entitlement policies set by IT.

Monitor software utilization
HP Configuration Management Application Usage Manager monitors the utilization of every application on all of your desktops, notebooks and servers. With direct visibility into the location, frequency, version status and trends of software use, your IT organization can reduce costs and mitigate risks.

Secure software from vulnerabilities
HP Configuration Management Patch Manager provides full lifecycle management of patches, service packs and hot fixes, including discovery, download and collection, testing, conflict analysis and vulnerability assessment, targeting, deployment and continuous enforcement. By automating patch management, deployment time is decreased from months to days, thereby reducing the risk of security vulnerabilities.

Comprehensive reporting
HP Configuration Management software brings the reporting elements from every tool in the solution together for comprehensive and centralized operational reporting.

You can also extend reports to include your own data resources. Hundreds of reports and views are available out of the box with flexible customization. Depending upon business needs, you can drill down for greater levels of detail in problem areas, or get an executive dashboard for high-level IT operations status.

Centralized policy management for compliance
The Enterprise Manager administrative console is a web-based console for centralized policy administration. Administrators can manage multiple directory services from a single console and quickly identify directory objects, including policy assignment for each managed object. Enterprise Manager provides enhanced security with role-based administration and access controls by only defining access rights within the directory. In addition, all policy changes can be tracked in a complete audit trail to document what policy changes were made, at what time and by whom, for compliance purposes. With centralized and continuous management, entitlement policies are enforced to reduce the risk of unauthorized access to systems. In addition, software is maintained and deployed via an electronic Definitive Media Library according to ITIL best practices to ensure its integrity.

Part of a closed-loop change management solution
HP Configuration Management software is tightly integrated with HP ServiceCenter and HP AssetCenter to form a robust solution for closed-loop change and asset management. The change process, from request through deployment and verification, can be managed through the HP ServiceCenter console and synchronized with HP Configuration Management software. In addition, tight integration with HP AssetCenter provides closed-loop asset management and enables software license compliance. By providing an automated, closed-loop solution, IT can deliver new services faster, more reliably and with greater efficiency.

In summary, the HP Configuration and Patch Manager provides the following features:

• Automated deployment: Efficiently, reliably and quickly deploy software changes across the largest number of devices, from hundreds to hundreds of thousands to reduce management costs, time to market and risk.

• Security and compliance enforcement: Define and centrally manage the policies governing software configurations across the enterprise, automate patch management, maintain an audit trail for compliance and enable compliance through policy enforcement and software distribution based upon an electronic Definitive Media Library.

• Continuous lifecycle management*:* Utilize a common automation tool to manage heterogeneous and distributed servers, desktops and notebooks for their entire lifecycle - discovery, provisioning and deployment, ongoing management and updates, and software removal and retirement.

• Windows Vista support: Reduce the time, cost and risk of Windows Vista migrations, including Windows Vista readiness evaluation and reporting, conflict analysis, and automated migration and deployment.

• Closed-loop change management*:* Integrates with HP ServiceCenter and HP AssetCenter to automate the change process for IT efficiency and acceleration of service delivery.

4.2.3.3.2. ProLiant Essentials Vulnerability and Patch Management Pack

For more specific security patch management, the ProLiant Essentials Vulnerability and Patch Management Pack integrates comprehensive vulnerability assessment and advanced patch management functions with HP Systems Insight Manager. It identifies and resolves security vulnerabilities quickly, efficiently, and reliably. The pack can be used independently or integrated with a broader patch management solution like the HP Configuration Manager Patch Manager.

Features of the ProLiant Essentials Vulnerability and Patch Management Pack include:

• Combined vulnerability assessment and patch management: A single tool seamlessly combines the assessment and the remediation of vulnerabilities, reducing the operational complexity that arises from managing separate tools.

• Integration with HP Systems Insight Manager: Integration enables use of existing functionality (such as discovery, identification, scheduling, role-based security, notification, and group-based actions) to eliminate the need for users to recreate tasks in multiple tools for vulnerability assessment and patch management.

• Comprehensive vulnerability assessment: Coverage of vulnerabilities reported in all leading vulnerability databases ensures comprehensive assessment. Powered by Harris STAT Scanner (the only scanner with Common Criteria Certification, an internationally accepted security qualification), the assessment identifies vulnerabilities reported in the Common Vulnerabilities and Exposures (CVE) list, the Federal Computer Incident Response Center (FedCIRC) vulnerability catalog, the SANS Top 20 Internet Security Vulnerabilities list, the CERT/CC advisories list, and the U.S. Department of Energy Computer Incident Advisory Capability (CIAC) bulletins.

• Acquisition, deployment, and enforcement of patches: The pack automatically collects new vulnerability updates and patches directly from vendor sources, such as a vendor's web-based patch depository. Patch manifests, which break down each patch into its component parts, are created automatically.
• Centralized management: Schedulable patch deployment, patented differencing (differences between actual and expected configurations), and checkpoint restarts (resuming processes at checkpoints due to interruptions) ensure that patches are deployed with minimal impact on network resources and allow patches to be managed from a central point.
• Unique desired-state management: The system automatically and continuously ensures that patches remain applied in their proper state. If patches are corrupted in any way, they are automatically reinstalled to bring the system to the desired patch level.
• Server lifetime coverage: The license provides coverage for the lifetime of the server for software upgrades and vulnerability updates.

4.2.3.3.3. Microsoft Security Patch Management Tools

Microsoft Corporation provides several tools to help with security patch management:

• The Microsoft Baseline Security Analyzer (MBSA)
• Automatic Updates (AU)
• Windows Server Update Services (WSUS)
• Systems Management Server 2003 (SMS 2003) and the Systems Management Server 2003 Inventory Tool for Microsoft Updates

The MBSA tool can perform a general security analysis scan on Microsoft Windows NT 4.0 and later versions of Microsoft Windows systems (the latest MBSA version 2.1is fully compatible with Windows Vista). MBSA can also scan the following Microsoft applications: Exchange Server, SQL Server, Microsoft Office and ISA Server. Besides a security analysis scan, MBSA also provides patch-scanning functionality. In addition, MBSA can be integrated with WSUS. This means that MBSA can check the enterprise WSUS server for security updates instead of checking the Microsoft Corporation web site.

AU is the client-side patching and update engine that is integrated in the Windows client and server OS platforms. AU can leverage either WSUS or the Microsoft Update web service to obtain the latest security patches.

The Microsoft Update, Windows Update and Microsoft Office Update web services allow Windows users to easily download and install the latest Microsoft OS and application patches.

The Windows Server Update Services (WSUS) give enterprise administrators the ability to provide Microsoft Update-based security patch services to their users and systems in a controlled and secure manner. WSUS can be used to set up an enterprise update server from which internal Microsoft Windows clients and servers can download the latest patches.

Systems Management Server 2003 and the Systems Management Server (SMS) 2003 Inventory Tool for Microsoft Updates (ITMU) are Microsoft Corporation's most advanced security patch management tools. The ITMU can determine security patch status and generate reports on patch status. ITMU integrates with SMS 2003 for distributing and installing patches.

4.2.3.3.4. HP-UX Security Patch Management Tools

Security Patch Check (SPC) automatically downloads the latest security bulletin catalog and analyzes a system or depot. It then generates a report of applicable security bulletins and identifies required actions, including required patches, updates, software removals, and manual actions. To perform regular analyses, administrators can run it as part of an automated process (for example cron on UNIX and UNIX-like OSs) or set up automatic runs via Bastille, an open-source lockdown tool.

The HP IT Resource Center (ITRC) patch download page can be used in conjunction with SPC. It performs dependency analyses on the requested patches, ensuring the administrator has all needed patches.

4.2.3.3.5. Linux Security Patch Management Tools

Most, if not all, Linux distributors provide highly granular security patches on nearly a daily basis. Patching is part of a well balanced change-management process which includes tools that list applicable patches and dependencies such that they can be reviewed by application owners and system stakeholders before they are applied. Linux systems have a wide variety of these tools at their disposal, including APT, Yum, Yet Another Setup Tool (YaST), up2date, and HP Radia Patch Manager.

4.2.3.3.6. Summary of Patch Management

Security patch management is paramount to maintaining a proactive stance against threats and vulnerabilities. A number of tools available from HP help organizations to manage security independently and effectively across various OSs and platforms. For more effective business continuity and enterprise-level protection, security patch management tools can be tied into the larger IT management function via integration with a general patch management tool. In addition, HP's Proactive Security Management Program identifies and prioritizes critical patches. It conducts system audits to ensure compliance and calculates business risks associated with newly identified vulnerabilities. HP can provide the right level of solution for any business requirement.

4.2.4. Lockdown and Hardening

Traditionally, systems have shipped with all of their features and capabilities turned on. If these capabilities are not in use, the system can be in a more vulnerable state than necessary. Lockdown and hardening consists of turning off unneeded services and features, configuring the remaining features and services to restrict data flow to only those that need it, and finally configuring applications through secure protocols to be more resilient to attack.

It might seem that locking down or hardening a system is a one-time task completed during initial installation; however, configurations can be altered-accidentally or maliciously. For this reason, hardening must be audited and regularly verified as part of the compliance monitoring component of an effective proactive security management program.

Hardening and lockdown methods are platform specific. HP provides capabilities to lockdown systems running on HP-UX, Linux, and Microsoft platforms.

4.2.4.1. HP-UX Lockdown and Hardening

HP-UX Bastille is an open source, security-hardening tool supported by HP for use on the HP-UX OS. It is the first comprehensive lockdown tool to provide an intuitive, educational, wizard-style interface, making it is easy to use for non-experts. HP-UX Bastille allows inexperienced and experienced security administrators alike to quickly make appropriate security decisions and tradeoffs.

Interactive elements in the user interface educate the system administrator about security issues. Bastille saves users time and pain with its supported and tested configuration changes. Furthermore, Bastille's "ratchet" lockdown approach ensures that users do not accidentally "loosen" their system. The tool also provides a revert feature as a safety net to quickly remove the Bastille security configuration if needed.

HP-UX Bastille can operate:
- Interactively with the wizard interface to harden the local host or to create a generic profile for use on multiple hosts
- At installation time via the install-time Security Ignite-UX interface

HP-UX Bastille performs a number of specific lockdown actions. These include:

- Removing risks associated with unused features by configuring system daemons, kernel, OS settings, network, and software
- Lowering patch urgency for disabled products by disabling unneeded services, such as echo and finger
- Providing additional security layers for Internet services such as web and Domain Name Service (DNS) by creating chroot "jails"
- Assisting patch currency by configuring Security Patch Check (SPC) to run automatically
- Dramatically reducing the system's network exposure by configuring a simple, comprehensive, deny-all inbound IPFilter firewall

For more information about HP-UX Bastille, see **www.hp.com/products1/unix/operating/security/index.html**.

4.2.4.2. Linux Lockdown and Hardening

Many Linux OS distributions also contain lockdown applications, including Bastille. Bastille examines the system and walks the user through a system-hardening process. It will not make assumptions or modify the system without getting approval for each step. Bastille can be used to iteratively harden a system, while LogCheck, PortSentry, Tripwire, and AIDE can be used independently as variable guides to determine if a system is being (or has been) compromised. For more information, see www.bastille- linux.org.

Tiger Analytical Research Assistant (TARA), which has been extended to comply with HP's own IT Security Controlled Host Requirements, can scrub the system for anomalies that might compromise the system's integrity (such as anything that might make the system fail an audit). HP development engineers have extended TARA to encompass the IT Security Linux Controlled Host Requirements for Internet facing network environments. With these and other tools, and an appropriate security review process, system administrators are able to significantly increase and maintain the security of newly deployed systems.

4.2.4.3. Microsoft Windows Lockdown and Hardening

Microsoft provides a lot of guidance for locking down its OS platforms and applications. Good examples are the Windows XP Security Guide, Windows Vista Security Guide and Windows Server 2003 Security Guide.

In Windows Server Active Directory (AD) domain environments, Windows platforms can be easily locked down by following the guidance that is given in the above security guides, and by using the Security Configuration Wizard (SCW - a built-in security lockdown tool for Windows servers and clients) and the Group Policy Object (GPO) security and configuration settings.

In addition HP provides the following Windows OS and Microsoft application hardening services and solutions.

HP Services provides hardening services for Microsoft clients and the following server roles bundled with the Windows Server OS:

- Domain controller
- Directory server (Active Directory (AD) or Active Directory Application Mode (ADAM))
- Dynamic Host Configuration Protocol (DHCP) server
- Domain Name System (DNS) server
- Windows Internet Naming Service (WINS) server
- File server
- Print server
- Internet Information Server (IIS) (Microsoft's application server)
- Internet Authentication Service (IAS) server (Microsoft RADIUS server)
- Certificate server (for Public Key Infrastructure services)
- Network Access Protection (NAP) (Microsoft's Network Admission Control (NAC) solution)

Specific server roles (not bundled with the Windows Server OS and part of dedicated Microsoft software offerings):

- Exchange Server: Messaging server
- Office Communications Server (OCS): Real-time collaboration server
- SharePoint Portal Server: Web portal server
- SQL Server - Database server
- Identity Lifecycle Manager (ILM): Identity management and provisioning server
- System Center: Management server
- BizTalk Server: Business integration and process management server

HP also offers exclusive Microsoft Windows NT, Microsoft Windows 2000, Microsoft Windows XP, Windows Vista and Microsoft Windows Server 2003 security solutions, called HP ProtectTools for Microsoft Products. These solutions can, for example, replace the password hashing algorithms supplied in Microsoft Windows with customer-specific algorithms that make brute force or dictionary password hacking much more difficult.

In the context of this discussion on Windows hardening, the following are relevant HP ProtectTools for Microsoft Products solutions:

• HP ProtectTools Authentication Services: This product provides a number of features that enhance the standard Microsoft authentication process. The central feature is enhanced password management, achieved by implementing a CESG-approved password hashing and password generation system. Each organization is provided with special CESG seed values to ensure each organization's system is unique. Where government algorithms are not applicable, alternative commercial algorithms are used. The product also manages change of administration passwords, provides last successful and unsuccessful login information, and can be configured for multiple login denial and timed auto-logout. The use of a unique password hashing mechanism for systems prevents access from unauthorized systems even when a valid username and password are used.

• HP ProtectTools Windows Mobile: Most organizations see mobile commuting as the next big opportunity for achieving cost reductions while increasing business efficiency. Security has been the major concern preventing the take-up of this technology. HP has the capability to secure remote connections and protect the data held on mobile devices such as laptops and PDAs. HP ProtectTools Windows Mobile ensures that proper authentication is undertaken, that all data is deleted if the system is lost or stolen and the password is incorrectly entered a predetermined number of times, and ensures that PDAs can only link with known and authorized PCs.

5. Proactive Security Management Summary

It is often joked that the most secure computer is one that is in a guarded, locked room…and is also turned off. The point of the joke is that there is no such thing as 100% security and the most secure system is one that is not useful. The reality is that there is a set of trade-offs or variables to manage such as costs, asset values, security technologies, and people. Proactive Security Management is the science of managing those variables with people, processes and technology to support an organization's goals, and do so while maintaining an acceptable level of risk. The environment for our IT infrastructures includes an ever-changing state of threats, an evolving set of vulnerabilities, and the basic, human-nature condition that if something has value, then there's at least one person who might try to take it.

To be certain, Security Management has matured far beyond simply keeping the bad guys out or presenting a single console to corroborate security point tools. In order to achieve its stated goals, security management must: (1) Manage the protection of data, applications, systems, and networks, both proactively and reactively; (2) respond to changes in business and organizational models as well as the changing threat environment; (3) integrate with IT infrastructure management and operations; and (4) all the while, maintain a level of security and operational risk that is pre-defined by that organization.

For further information about Proactive Security Management products and solutions from HP, please see the following URL locations:

Table 2-2
HP proactive security management offering summary

Proactive Security Management Solutions	www.hp.com/go/security/proactive
Proactive Security Management Services	www.hp.com/go/security Click HP security services link
Security Assessment Services	www.hp.com/go/security Click HP security services link
SMB Security Services	www.hp.com/sbso/services/security_vulnerability.html
Security Health Check Services	www.hp.com/go/security Click HP security services link
Global Security Operations Centers	www.hp.com/go/security Click HP security services link
Security and HP-UX 11	www.hp.com/go/hpux11isecurity

Chapter 3
Identity Management

"The increasingly distributed nature of corporate networks, the proliferation of web-based applications, increased security awareness, and government regulations such as Sarbanes-Oxley and HIPAA have contributed to making identity management a necessity for virtually every business."

-Roberta Witty, Research Director, Gartner, Inc.

Figure 3-1
Identity management

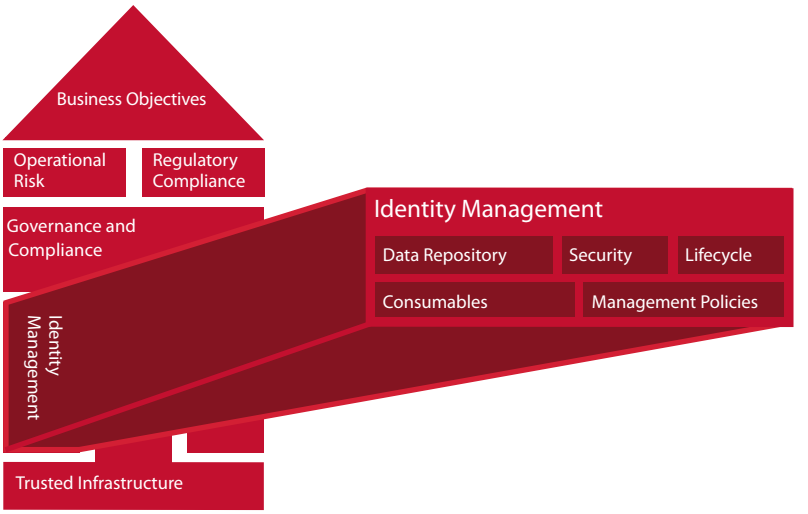

Identity management is one of the three key security areas in which HP is innovating. Within the HP security framework, identity management provides a set of processes and tools that allow administrators to manage large populations of users, applications, and systems quickly and easily. In addition, business policies, regulatory compliance, and risk factors shape the policies and practices that direct identity management.

This chapter begins by providing the definition and purpose of identity management. Next, it presents the identity management components, and key elements of identity management solutions. The final section of the chapter discusses the specific identity management capabilities that HP delivers.

"Organizations are using identity and access management to reduce administrative costs, increase user productivity, tighten security and make systems compliant with policies and regulations."
-Carol Baroudi, Aberdeen Group, "Identity and Access Management" Report, March 2007

1. Definition

Identity management is the set of principles, processes, tools, and social contracts surrounding the creation, maintenance, and use of digital identities for people, systems, devices and services. It enables secure access to a set of systems and applications. Identity management solutions and infrastructures include data repositories, security services, lifecycle management services, consumable value, and management components. Identity management has strong links to security, trust, and privacy management. It also delivers components of risk management.

Traditionally, identity management has been a core component of system security environments. It is used for maintaining account information and controlling access to a system or limited set of applications. Control is usually the primary focus of identity management. For example, an administrator issues accounts to restrict and monitor access to resources. More recently, however, identity management has also become a key enabler of electronic business.

2. Purpose

Identity management combines processes and technologies to secure and manage access to an organization's resources and assets. In addition, it identifies every user (even anonymous ones), application, service or device throughout and across organizations, and over time. Identity management provides flexible authentication, access control, and auditing while respecting privacy and regulatory controls. Identity management systems are fundamental to establishing accountability in business relationships, customizing the user experience, protecting privacy, and adhering to regulations.

The following list provides examples of the primary goals that drive organizations to implement identity management solutions:

- Reduce total cost of ownership (TCO) for all systems, including the costs of administration, help desk, and technical support
- Reduce management overhead
- Provide competitive advantage by enabling automation and streamlining optimization of business processes
- Improve customer and employee service
- Support the maintenance, the control and confidentiality of customer, supplier, and employee data
- Reduce the time for employees, partners, customers, services, devices and others (e.g. contingent and emergency workers) to gain access to required organizational resources
- Reduce the risk of using incorrect information for business processes
- Reduce the risk of employees, partners, customers, services, devices and others retaining access to organizational resources after their relationship with the organization has changed (e.g. promotion of an employee) or ended (customer ends service contract)
- Support legal and compliance initiatives related to employee and customer data, for example, the Health Insurance Portability and Accountability Act (HIPAA), Sarbanes-Oxley, the EU Directive on Data Protection, the Basel II Accord, and the Canadian Privacy Act

In short, the purpose of identity management is to provide organizations the following key benefits:

- Enhanced enterprise agility and productivity
- Improved end-user convenience
- Increased IT management efficiency, including cost reduction
- Effective regulatory compliance

Figure 3-2
Identity content

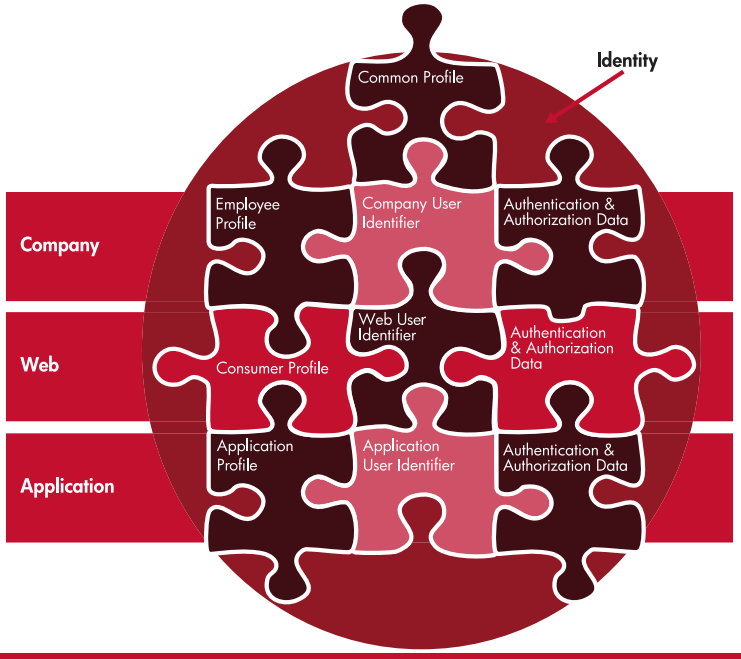

3. What is a Digital Identity?

"Digital identity is one of the fundamental building blocks for the next generation of information systems."
- Tony Scott, CTO, General Motors

Identity is a complicated concept with many nuances that range from the philosophical to the practical. In the context of identity management, however, the identity of an individual is the set of information known about that person. In the digital world, a person's identity is typically referred to as a digital identity. Different contexts, roles, affiliations, and application environments can require different levels of assurance and digital identities. Therefore, a person can have multiple digital identities. These potential multiple identities are often referred to as personalities, profiles, guises, avatars and other similar terms.

Although digital identities are predominantly associated with humans, they are increasingly being associated with non-human entities (services, systems, and devices) that could act on behalf of people. Specific examples include trusted platforms, next-generation mobile phones, identity-capable platforms and Digital Rights Management (DRM)-based devices.

Figure 3-2 illustrates the content of a digital identity. Identity consists of a person's unique profile data, identifier data, and authentication and authorization data. Each content piece can be linked to different contexts (company, web, and application) and the person's role in that context. For example, a person's identity can be made up of a set of names, addresses, driver's licenses, passports, field of employment, etc. This information can be used for identification, authentication and authorization purposes, for example:

• A name can be used as an identifier - it allows us to refer to the identity without enumerating all of the items.

• A passport which can be used as an authenticator - they are issued by a relevant authority and allow us determine the legitimacy of someone's claim to the identity.

In different contexts, different unique identifiers can be used. For example, in the above example the driving license could be a relevant unique identifier for interacting with the Department of Motor Vehicles; and name-surname-address is the unique identifier for the post office or a delivery service, and so on. It is also important to consider the information and potential identifiers that can be derived or aggregated from such data. For example, in many countries a driver's license, while functioning as a privilege, may also contain other information such as a name, address and a photo allowing it to serve as an identifier or authenticator. This information can be used to link to other systems, going well beyond the function of proving the privilege to drive.

Often data associated with an identity is used improperly. For example, the use of a birth certificate as an authenticator represents a particularly poor choice, for there is generally nothing about the birth certificate that allows an individual to be correlated to the claims on the certificate. A better choice would be to use a passport holding an individual's photo *ID*, or a smart card storing an individual's fingerprint, or better still, multiple authenticators at the same time. Ultimately each of these documents is derived from an initial claim of identity which is often established by presenting a birth certificate.

This is especially important during the identity verification phase, where an organization first establishes a relationship with an employee, partner, customer or otherwise. This process is a critical step in ensuring that future actions and processes are maintained to an acceptable level of protection.

A high value is placed on authenticators that have gone through a rigorous vetting process prior to issuance. For example, despite the intention, and government guidance, not to use the U.S. Government's Social Security Number (SSN) as a unique identifier for anything other than financial institutions, many organizations do so. Because of the ubiquitous and generally accepted use of this identifier, its value as a unique personal identifier has decreased as identity thieves look to take advantage of that ubiquity. The value of an identifier is directly related to how it is connected to an individual; the stronger the linkage the more valuable the identifier.

Metadata (information about data) qualify all identity data, and an organization's policies for identity, authentication, authorization, and privacy protection define the metadata requirements. Policies are defined by an organization's IT and business decision makers - they are aligned with corporate governance rules, regulatory restrictions, and contractual obligations specific to the organization's operating environment.

Identity information and related policies can change over time. This means that identity management not only deals with static information but also copes with changes to identity data. The same is true for security policy management.

Multiple views can exist on an entity's identity information. Each view defines a digital identity that is valid and appropriate based on the context or purpose. Using multiple views within and across multiple contexts enables interactions and transactions. Examples of different views and contexts are illustrated in Figure 3-3.

Different stakeholders can disclose, access, and use digital identities in one or more contexts, including personal, social, e-commerce, enterprise, and government. The process occurs through a variety of means including personal appliances, enterprise systems, and web services.

From an identity subject's point of view, there are multiple perceptions of their identity information:

- **Me Me** is the part of identity information that the subject is aware of and directly controls. An example is personal address information stored and maintained in an organization's white pages directory. It can also include personal or private information - such as a credit card number or SSN - that an individual carefully protects and reveals only in particular circumstances.
- **Known Me** is the part of identity information that the subject is aware of and indirectly controls. An example is an individual's revenue data and associated tax levels that are stored in the tax department's database. Even though an individual provides the revenue data to the tax department, he or she doesn't have direct control of the content in the database.
- **Unknown Me** is the part of identity information that the subject is not aware of and cannot control. Other stakeholders, which may be known by the subject, can control this information. Examples include Certification Authorities (CAs), authorized e-commerce sites, Trusted Third Parties (TTPs), and unknown third parties (for example, credit rating agencies and identity thieves).

Figure 3-3
Identity views and contexts (Subject's perspective)

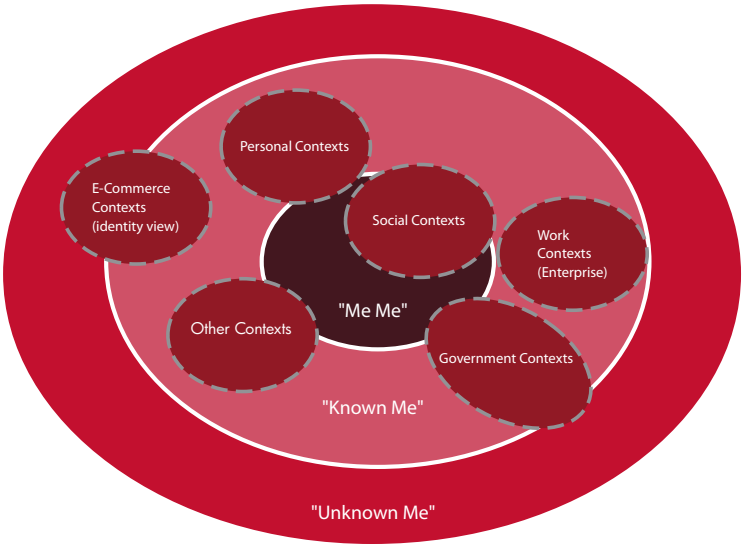

Figure 3-4
Identity management components

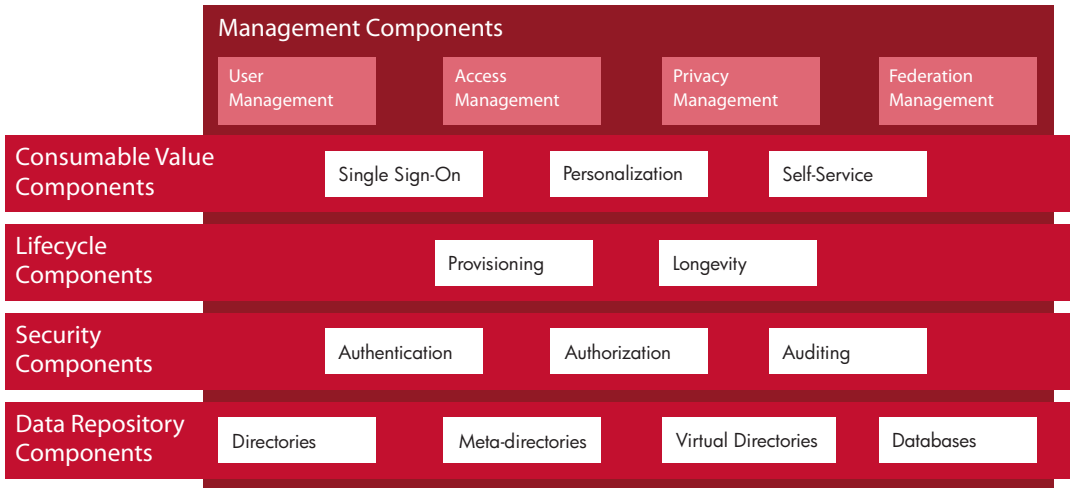

4. Identity Management Components

Identity management solutions are modular and composed of multiple service and system components. This section outlines the components of a typical identity management solution, as illustrated in Figure 3-4.

Components of identity management solutions exist at different maturity stages. Components like authentication and directories are very mature and are considered consolidated technologies.

Provisioning, authorization, federation and single sign-on (SSO) components are rapidly consolidating. Others, such as the privacy management component, are still in a definition and research stage.

4.1. Data Repository Components

Identity repositories deal with the representation, storage, and management of identity and profiling information. They provide standard Application Program Interfaces (APIs) and protocols for information access. Data repositories are often implemented as a Lightweight Directory Access Protocol (LDAP)-accessible directory, meta-directory, or virtual directory. Other repositories that are used in the context of identity management solutions are databases and XML-formatted files. Policy information, which governs access to and use of information in the repository, is generally managed and stored in these repositories as well.

4.2. Security Components

Authentication providers, sometimes referred to as identity providers, are responsible for performing the primary authentication that links an individual to a given identity. The authentication provider produces an authenticator - a token allowing other components to recognize that primary authentication has been performed.

Primary authentication techniques are generally considered in terms of:

- "Something you know", such as a password, PIN or the answer to an identifying question (e.g. mother's maiden name).
- "Something you have", such as a mobile phone, credit card, an X.509 public-key infrastructure (PKI) certificate, smart card, or other hardware security token.
- "Something you are", such as a fingerprint, a retinal scan, or other biometrics.
- "Somewhere you are", such as being within a given geographical location, or within range of a known beacon.

Each identity may be associated with multiple authentication providers. In addition, the mechanisms employed by each provider may be of different strengths. To accept the claim of a given identity, some application contexts may require a minimum level of strength.

Authorization providers enforce access control when an entity accesses an IT resource. Authorization providers allow applications to make authorization and other policy decisions based on privilege and policy information stored in the data repository. An authorization provider can support simple access control management at the operating system (OS) level. It can also support finer-grained role- and/or rule-based controls at the application and service levels.

Auditing providers supply mechanisms to track how identity information in the data repositories is created, modified, and used. They are an essential enabler of forensic analysis, which helps determine who circumvented policy controls and how the controls were evaded.

4.3. Lifecycle Components

Provisioning is the automation of all the procedures and tools used to manage the lifecycle of an identity. Provisioning procedures include:

- Creating the identity including an identifier
- Linking to authentication providers
- Setting and changing attributes and privileges
- Decommissioning an identity

In large systems, provisioning tools generally permit some form of self-service for creating and maintaining an identity. They frequently use a workflow or transactional system to verify data from an appropriate authority. The tools may also propagate data to affiliated systems that may not directly consume the repository.

Longevity tools create the historical record of an identity. These tools allow the examination of the evolution of an identity over time. Longevity is linked to the concept of attestation - the ability to determine which actors had access to which resources and in what timeframe (irrespective of whether they exercised access, which is a matter of auditing).

4.4. Consumable Value Components

Identity Management solutions can provide the following value-added services for the users or consumers of an identity management system: SSO, personalization and self-service.

SSO allows a user to perform a single primary authentication for access to the set of applications and systems in the identity management environment.

Personalization and preference management tools associate an identity with application-specific and generic information. These tools allow applications to tailor the user experience, streamline the user interface, and target information dissemination for a business.

Self-service enables users to self-register for access to business services and manage profile information without administrator intervention. It also allows users to manage their proper authentication credentials; for example, assigning passwords, resetting passwords, and requesting X.509 PKI certificates. Self-service reduces IT operation costs, improves customer service, and improves information consistency and accuracy.

4.5. Management Components

User management provides IT administrators with a centralized infrastructure for managing user profile and preference information. User management enables organizations to decrease overall IT costs through directory optimization and profile synchronization. User management tools provide user self-service capabilities and enhance the value of an organization's existing IT investments.

Access management provides IT administrators with a centralized infrastructure for managing user authentication and authorization. An access control management service increases security, reduces complexity, and reduces overall IT costs by automating access policies for employees, customers, and partners.

Privacy management assures that identity management solutions respect privacy and data protection policies as defined in company, industry, and governmental regulations, while storing, accessing, processing and disclosing personal data.

Federation management establishes trusted relationships between distributed identity providers. This often involves sharing web service endpoints, X.509 PKI certificates, and supported/desired authentication mechanisms.

4.6. The Effect of Policies on Management Components

Policy controls govern and drive management components. Policies may cause events to be audited or an identity subject to be notified when information is accessed. The following policies are typically involved in an identity management solution:

- Identity policies control the format and lifetime of an identity and its attributes.
- Authentication policies control the characteristics and quality requirements of authentication credentials.
- Authorization policies determine how resources can be accessed.
- Privacy policies govern how identity information may be accessed, processed or disclosed.
- Provisioning policies determine what resources are allocated to which identities and how the resources are allocated and de-allocated.

5. Key Elements of Identity Management Solutions

There are many products and solutions available in the identity management market. They generally provide one or more of the identity management components and target different types of users and contexts, including e-commerce sites, service providers, enterprises, and government institutions. Key IT industry players are currently focusing on creating identity management suites that provide all of the components shown in Figure 3-4.

There is a considerable amount of overlap between the different solution categories available on the market. A good example is meta-directories and provisioning solutions. The role of meta-directories has gradually shifted from pure data synchronization (a repository function) to lifecycle component functions for the creation of user entries (a provisioning function).

Identity management solutions also involve other stakeholders. These include authentication devices (smart cards, biometric devices, and authentication tokens); anonymity services; and the standards outlined in the next section.

The quality of identity management products and solutions depends on how successfully they handle a number of factors. Among other things, these factors include keeping identity information in a consistent and up-to-date state, satisfying related management policies and legal requirements, preserving privacy and trust, and ensuring that security requirements are fulfilled. The key elements to consider in an identity management solution include:

- Adherence to identity management standards
- Types of deployment models
- Means of addressing complexity and competing demands
- Methods of safe digital identity management
- Level of product interoperability

Organizations should also consider the following identity management trends when building an identity management solution:

- Identity services
- Business-driven identity management
- Identity-capable platforms and device-based identity management

We will explain these key elements and trends in more detail in the following sections.

5.1. Identity Management Standards

Standards provide a common set of protocols, semantics, and processing rules that allow the various components of an identity management solution to interoperate. Table 3-1 provides an overview of the most important current and emerging standards used in identity management architectures and solutions.

Recently there has been an increase in the standards that are proposed by one or a few companies that are very active in the identity management sector (as opposed to the traditional consortium-driven identity management standards). Good examples are the OpenId initiative and the Microsoft CardSpace initiative (that leverages WS-*).

Table 3-1
Relevant standards for identity management architectures

Type	Relevant Standards	More Information
Identity Repositories	Lightweight Directory Access Protocol (LDAP)	www.ietf.org
	ISO/ITU x.500	www.itu.int
Privacy Standards	Platform for Privacy Preferences (P3P)	www.w3.org/P3P
	Enterprise Privacy Authorization Language (EPAL)	www.zurich.ibm.com/security/enterprise-privacy/epal
Access Control	eXtensible Access Control Markup Language (XACML)	www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml
Provisioning Authorities	Service Provisioning Markup Language (SPML)	www.oasis-open.org/committees/tc_home.php?wg_abbrev=provision
Federation	Security Assertion Markup Language (SAML)	www.oasis-open.org/committees/tc_home.php?wg_abbrev=security
	Liberty Alliance Standards	www.projectliberty.org
User-centric Identity Management	OpenID	www.openid.net
	Microsoft CardSpace	www.microsoft.com/net/cardspace.asp
	Higgins	www.eclipse.org/higgins/
Supporting Standards	WS-Security	www.oasis-open.org/committees/tc_home.php?wg_abbrev=wss
	WS-* (Roadmap)	http://msdn2.microsoft.com/en.us/library/ms977312.aspx

5.2. Deployment Models

Identity management systems are primarily deployed in one of three models: *silos*, *walled gardens*, or *federations*.

• **Silos** are the predominant model on the Internet today. In this model, the identity management environment is established and operated by a single entity for a fixed user and resource community. A good example is a Microsoft Windows domain governed by a set of predefined administrators and domain controller (DC) servers.

• **Walled gardens** represent a closed community of organizations. A single identity management system serves the common user community of a collection of businesses. Most frequently, this occurs in business-to-business exchanges, and specific rules govern the entity operating the identity management system. A good example is the Identrus PKI, which brings together individual bank-level PKIs into a closed banking-community PKI.

• **Federations** and federated identity management environments are emerging deployment models. They include systems like Liberty Alliance Project-based federations, federation systems built on the Web Services Security (WS-Security) and WS-Federation (WS-Fed) standards, and relatively recent user-centric frameworks such as the OpenId and Microsoft CardSpace initiatives.

The central difference between federated identity management systems and walled gardens is that a single entity operates a walled garden. By contrast, federated systems support multiple identity providers and a distributed and partitioned store for identity information. Clear operating rules govern the various participants in a federation - both the operators of components and the operators of services rely on the information provided by the identity management system. Most systems exhibit strong end-user controls over the dissemination of identity information to members of the federation.

5.3. Complexity and Competing Demands

The current identity management landscape is very complex because of the multiple interests, perspectives, concerns, and technologies that are involved. Identity management is important in different scenarios, including enterprise, e-commerce, social networking and government. It supports business processes and services, and it enables digital interactions and transactions.

There are competing demands on what identity management should provide, differing concerns about its focus, and conflicting interests. Examples of competing demands include enterprise focus versus consumer focus, mobility versus centralization, legislation versus self-regulation, subjects' control versus organizations' control, and privacy versus free market.

5.3.1. Numerous Stakeholders

Demands are dictated by various stakeholders, which can include enterprises, e-commerce sites, service providers, government agencies, and identity subjects (consumers). Stakeholders have different objectives and priorities when dealing with the management of digital identities:

• Enterprises are driven by their business objectives and needs. They manage large sets of identity data to enable their businesses, rationalize their assets, simplify interactions with partners and customers, ensure regulatory compliance, and meet contractual obligations. Identity data also helps enterprises manage the information lifecycle of their workforce and manage access to enterprise resources.
• E-commerce sites and service providers manage consumers' identity information to achieve a variety of goals, such as increasing sales, understanding customers' needs, customizing services, and selling information to third parties.
• Government agencies are concerned with the control and protection of their citizens' personal information. They also seek strong and undeniable authentication mechanisms and the automation and rationalization of their services via the Internet.
• Consumers have different concerns and needs depending on the role they play. They are in the middle (or, depending on the point of view, the source) of most of the competing demands previously noted. As employees or consumers, they want to access and use services in the simplest and most efficient way. Private citizens' needs and concerns might include privacy, distrust of institutions, and the accountability of the involved parties.

This variety of interests and concerns, along with emerging technologies, increases the complexity of identity management. All of these aspects influence each other, via a spiral of potentially conflicting requirements. For example, new legislation addresses citizens' needs for privacy; however, it constrains how enterprises, service providers, and e-commerce sites process personal information.

5.3.2. Multiple Domains

Multiple domains can also increase the complexity of identity management. Business tasks, digital interactions, and digital transactions can span multiple domains. In an e-commerce context, for example, a digital transaction might require the involvement of identity e-commerce sites and the exchange of identity information among these sites. This exchange has strong implications for managing trust, privacy, authentication, authorization, and accountability. Business-to-business interactions and transactions within supply-chain communities face similar challenges as a result of multiple domains.

5.3.3. Fragmented Implementation

Further complexity derives from the challenge of installing, configuring, administering, and integrating current identity management products. This is mainly due to the fragmentation of identity management components and the lack of interoperability and standards. This complexity creates frustration and delays the adoption of identity management solutions in the IT environment.

5.4. Safe Digital Identity Management

Identity management systems bring great value to the digital world. Federated identity environments, in particular, hold promise for widespread deployment. As the distinction between real-world identity and digital identity blurs, however, methods of safe digital identity management need to be considered:

- Authenticity of identity or Identity Assurance. How is the accuracy and validity of identity information measured and determined? What trust services or identity proofing processes are necessary to generate confidence in information in the identity management system?
- Longevity of information. Do identity management systems adequately track changes to identity information over time? Do they maintain the necessary artifacts to support historical investigations? How often does identity proofing or re-proofing need to take place? How is information disposed of and how often must the disposal occur?
- Privacy. Do identity management systems provide adequate controls to preserve individual privacy? Does the system provide adequate support for anonymity and multiple user-controlled personas?
- Identity theft. Do widespread identity management systems make it easier to perpetrate identity theft or identity fraud, and what can we do to minimize risk to the organization, its employees, partners, customers and others?
- Legal structures. What protections exist for the holder of the identity or the relying party? Do these protections go beyond contractual obligations when digital identity systems are used for interactions that are limited to the physical world today?

5.5. Product and Solution Interoperability Challenges

Most of the current identity management products and solutions rely on self-contained, stand-alone management and control tools. Little integration or interoperability is available with other management tools to deal with the management of security, trust, and privacy in an orchestrated way. To react to changes, identity management products and solutions need to evolve toward higher levels of interoperability, flexibility, and capability.

Particularly challenging are the interoperability issues in the federated identity management space, that are due to competing and overlapping proposals, such as the Liberty Alliance ID-FF, OpenId, Higgins, and CardSpace initiatives. A promising initiative in this context is the Concordia Initiative, which tries to create interoperability options and proposals among the above mentioned proposals and initiatives. More info can be found at www.projectconcordia.org/index.php/Main_Page.

6. Identity Management Trends

Organizations should not only focus on the key elements that were outlined above when building an identity management solution they must also consider the following identity management trends:

- Identity services
- Business-driven identity management
- Identity-capable platforms and device-based identity management

6.1. Identity Services

An emerging trend in enterprise identity management and federated identity management is the use of Identity Services, or services that leverage identity management components and solutions to provide reusable identity management capabilities across organizations. Examples of emerging identity services are:

- Authentication and credentialing services
- Single-sign-on services
- Authorization services
- Provisioning services
- Services for long-term archiving of identity information and related records
- Cryptographic services such as digital signing, time-stamping and encryption

Identity services reflect the increasing interest and shift of enterprise IT applications and solutions towards Service Oriented Architecture (SOA)-based and Web 2.0-based approaches. In the medium and longterm, identity services will significantly impact applications and services and the way organizations deal with identity management in general.

6.2. Business-driven Identity Management

Enterprises and other organizations are increasingly managing IT from a business perspective, to reduce costs, improve availability, tune capacity, optimize resource utilization, and to deal with risks and regulatory compliance.

In this context, the ITIL (IT Infrastructure Library) framework defines a set of best practices that are focused on aligning IT with business objectives. The ITIL best practices create a service-oriented culture, where there is an understanding that IT exists to support the business, that there is a commitment to deliver an agreed level of service, and that customers' satisfaction always comes first.

The ITIL core disciplines are centered on Service Support and Service Delivery. ITIL provides guidance in terms of Configuration Management, Change Management, Incident Management, Security Management (based on ISO/IEC 17799) and Audit Management.

Considering the increased importance identity management has in enterprises and the trend towards identity services, ITIL will play a key role in the definition of best practices and identity controls for identity management.

6.3. Identity-Capable Platforms and Device-based Identity Management

An important emerging area in identity management is device-based identity management. This is about the management of the identity of devices (smartphones, PDAs, laptops, etc.) in enterprise and federated contexts; using these devices to store - in a secure and trusted way - identity information; and enabling interactions and SSO between these devices, their users and other parties.

A considerable amount of work in this space has been done by HP Labs, together with HP businesses. Two research projects that are specifically worth noting are the HPL project on 1) Liberty Alliance identity-capable platforms and provisioning services; and the HPL project on 2) device-based identity management in enterprises.

Liberty Alliance Identity-Capable Platforms and Provisioning Services

Liberty Alliance's "Advanced Client Technologies" initiative aims at defining and specifying technologies that encompass a suite of advanced functionality in the areas of identity-capable devices, SSO, identity federation, service hosting, reporting and provisioning.

The key goal of this HPL R&D initiative is to enable users, via identity-capable platforms (ICP - such as laptops, smart phones, PDAs, etc.) to engage in federated, multiparty interactions and transactions (on the Internet or other networks) in a simplified and transparent way. At the same time these ICP devices store, process and potentially disclose identity tokens in a secure, private and policy-controlled way. Identity tokens are provisioned to ICP devices via provisioning services and by means of protocols specified by the Liberty Alliance.

Intel has been leading the definition of Identity-capable Devices specifications in the context of Liberty Alliance. Other players include Nokia, BT, Vodafone, Gemplus, NTT, Sun, Symlab and HP. A working group in Liberty Alliance (supported, among many, by Intel, BT, HP Software/HP Labs) is standardizing ICP properties and its capabilities along with the required back-end operational services. HP and HP Labs have developed and provided back-end service capabilities that are required to enable ICP.

Device-based Identity Management in Enterprises

The management of device identities is becoming a key requirement in enterprises where the identities of platforms and devices have become as important as the identities of humans to grant access to enterprise resources. In this context, access control systems need to understand which devices with what properties are being used to access resource, by whom and in which contexts. Trust in managed devices' identities is an important first step to enable this.

The separation between work, public and private aspects of a person's life is becoming more and more blurred. In particular, some devices are not only used for work-related matters but also for personal matters, such as accessing the web to retrieve information and make transactions, exchanging personal e-mails, making personal phone calls, storing and keeping track of personal information, calendars, etc.

From a user (individual) perspective, this trend further simplifies their day-to-day life by avoiding any unnecessary duplication of devices, tools and related efforts. From an enterprise perspective, the fact that devices are used by employees for a variety of purposes, introduces additional risks and threats, in particular about the integrity of these devices and their trustworthiness to access enterprise intranets and networked resources. An additional risk is that private devices (e.g. personal laptops, etc.) could also be used at work - with potential lower security and assurance levels (e.g. about installed software, patch control, local access control settings, etc.) than the ones mandated by the enterprise.

Current enterprise services, applications and information are mainly protected by traditional access control systems that usually only take into account human-based identities (via passwords, digital certificates, etc.) or (in more advanced situations) only human-based identities that are strongly bound to a given device. To have better control of accessed resources, it is becoming more and more important for enterprises also to explicitly identify what the identity of a device is, along with its properties - consider the identity of a device as a self-standing entity or the identity of a device as one of a group of known entities. Furthermore, trust and assurance is required about the authenticity and validity of a device's identity.

HP Labs R&D in this space has been focusing on modeling devices' identities, enabling their provisioning in heterogeneous enterprise systems, providing support for making and enforcing related access control decisions, and leveraging trusted computing capabilities of modern devices to deal with aspects of trust management.

7. Summary of Identity Management Concepts

From a technological and IT perspective, identity management is just one aspect of managing business solutions and the overall IT stack of networks, platforms, OSs, applications, middleware, and services. Identity management must be considered in a holistic way by including the management of security, trust, and privacy along with the management of policies, requirements, and changes. All of these aspects are interrelated and affect business solutions and the IT stack at different levels of abstraction.

The components within the identity management landscape are rapidly changing. Classic identity management components are consolidating. On the other hand, new components and standards are emerging, such as identity federations, identity for web services, and privacy management solutions. A good example is the linking of identity and network security solutions in the network access control (NAC) space. Where NAC begins by authenticating and validating the posture (or health) of a device attempting to connect to a network, the identity dimension ensures that the device is known, as well as the user of that device, before access is granted to specific parts of the network. Combining the various layers of device, network, identity and communication access controls into common policy models and known enforcement points allows an organization to deliver a more comprehensive defense in depth environment.

Identity management is also gaining importance. Future identity management solutions will play a more central role in the IT industry due to the pervasiveness and the increased presence of identity information in all components of the IT stack. Very important trends are also identity services, business-driven identity management, identity-capable platforms and device-based identity management.

8. HP Identity Management Products and Solutions

For HP, identity management is the ability to:

- Identify every user, application, and device throughout and across organizations over time
- Provide flexible authentication, access control, and auditing technologies, while respecting privacy and regulatory controls
- Bring management capabilities to individuals, small organizations, and large organizations via easy-to-use and understandable tools that cope with dynamic populations and business changes

HP's identity management vision is centered on the pervasiveness of identity management technologies and solutions:

- Identity management is about the management of user, application, service and device identities.
- Identity management is about the management of identities in different contexts: enterprises, small and medium businesses (SMBs), consumers, and the public sector.
- Identity management deals with the management of the entire lifecycle of identities and their attributes.

The following sections will define HP's identity management vision in more detail by exploring the different identity management building blocks. We will address identity repositories, security components, privacy management, identity lifecycle management, and federated identity management. In addition, we will discuss the Hewlett Packard National Identity Solution (HP NIS) as an example of how these building blocks are combined to create an integrated national identity management solution.

8.1. Identity Repositories

Directories are the most commonly used repositories for storing identity-related information. As mentioned previously, identity management solutions can incorporate other repositories, including SQL-rooted databases and XML-formatted files.

8.1.1. Types of Directory-based Identity Repositories

Different technological approaches exist for directory-based identity repositories: centralized enterprise directories, meta-directories, directory synchronization utilities, and virtual directories. Of these, only a centralized enterprise directory is a true identity repository. The other tools integrate and link different identity repositories:

• Enterprise directories can provide a single authoritative source for identity information throughout an enterprise. All users and directory-enabled applications rely on the identities stored in the enterprise directory. This is the ideal scenario. However, most enterprises cannot use this approach due to the presence of legacy service-and application-specific directories.

• Meta-directories provide a consolidated view of the identity data stored in different repositories. They also synchronize the data in the different repositories. A meta-directory resembles an advanced directory synchronization utility. Most meta-directory solutions come with workflow logic, and they overlap with many of today's identity provisioning solutions.

• Directory synchronization utilities are intelligent LDAP-based utilities that can synchronize identity data between different types of identity repositories - such as directories and databases.

• Virtual directories, unlike meta-directories, do not build a central repository - although there is usually some degree of caching capability inherent in the products to mitigate potential network performance and reliability issues. Instead, they rely largely on directory server or client functions to access the data stored in different directory sources. Virtual directories also allow for the creation of different application-specific views of directory data - e.g. a customer application view and an employee application view.

8.1.2. HP and Identity Repositories

HP considers directory identity repositories a mature market and uses a best-of-breed and customer preference approach. Table 3-2 gives a non-exhaustive overview of directory solutions.

HP offers the Red Hat/Netscape Directory Server for HP-UX 11i. The directory server is built into the foundation HP-UX operating environment. It provides the central database repository of user names and objects for system and application access.

Table 3-2
Directory solutions

Vendor	Directory Product	URL
Enterprise Directory		
HP	Red Hat/Netscape Directory Server	www.hp.com/go/redhat www.hp.com/go/hpux11isecurity
Novell	eDirectory	www.novell.com
Microsoft Corporation	Active Directory Active Directory Application Mode (ADAM) Active Directory Lightweight Directory Services (AD LDS)	www.microsoft.com
Critical Path, Inc.	Directory Server	www.cp.net
Oracle	Oracle Internet Directory	www.oracle.com
Sun Microsystems, Inc.	Sun Java System Directory Server	www.sun.com
Meta-directory		
Microsoft Corporation	Identity Lifecycle Manager 2007 (ILM 2007)	www.microsoft.com
Critical Path, Inc.	Meta-Directory Server	www.cp.net
Siemens	HiPath SIcurity DirX Identity	www.siemens.com
Directory Synchronization		
HP	LDAP Directory Synchronizer (Compaq LDSU)	www.hp.com
IBM Corporation	IBM Tivoli Directory Integrator	www.ibm.com
Virtual Directory		
Radiant Logic, Inc.	Radiant One Virtual Directory Server	www.radiantlogic.com

Table 3-3
Overview of authentication methods and the authentication factors they support

	Password or PIN	Smart Card or Token	Biometric Device	Biometric Device and Smart Card	Dial Back	Trusted Platform Module (TPM)
Knowledge	x	x		x		x
Possession		x		x		x
Biometric Data			x	x		
Location					x	

8.2. Security Components

This section discusses the triple-A components of an identity management solution: authentication, authorization, and auditing services. It provides details about the solutions HP offers in this space.

8.2.1. Authentication Technologies

Authentication is the process of verifying an entity's identity. Authorization credentials, which are uniquely linked to an entity, are typically used for verification. The security quality of authentication technologies largely depends on the following dynamics: the number of authentication factors, the authentication protocol, and the authentication method.

Multifactor authentication methods offer higher security quality than single-factor authentication methods. A good example of a multifactor authentication system is a smart card. It combines possession (of the card) and knowledge (of the card's PIN). Table 3-3 gives an overview of different authentication methods and the number of authentication factors they support.

Today's problems of identity theft and the misuse of identities and their attributes are accelerated by the ever-increasing amount of interconnected users, applications, and devices.

Many identity management solutions require the authentication infrastructure to support multiple authentication methods and protocols. This may be necessary when the environment supports internal and external users that access a variety of resources. When resources hold different values or contain sensitive information, different methods and protocols may also be necessary. Access to confidential information, for example, may require a stronger authentication method than access to information published on a corporate intranet. In some authentication infrastructures, this feature is known as graded authentication. This simply means that the resources and information a user is allowed to access vary depending on the strength of the authentication protocol and method.

8.2.1.1. Strong Authentication

Today's problems of identity theft and the misuse of identities and their attributes are accelerated by the ever-increasing amount of interconnected users, applications, and devices**.** To attain greater levels of authentication, identity management solutions require strong authentication. Over the last decade, strong authentication has been associated with both cryptography and multifactor-rooted authentication.

Cryptography-based authentication means that the authentication protocol includes cryptographic operations in the identity and credential verification process. Table 3-4 provides descriptions of popular strong user and device authentication solutions.

Biometric Authentication

Biometric authentication is a form of strong user authentication receiving substantial attention. A biometric is any measurable aspect of a human's physiology or behavior that can be reliably captured and used as a distinguishing identifier for that person within a defined population. Biometrics are an authentication mechanism that use the "what you are" about an individual to determine his or her identity.

Biometrics are a method of tying a claim of identity to an individual in a way that is not easy to spoof. Historically, passwords or PIN codes authenticated a person's claim of identity. Passwords have many security-related issues, including the ease with which they can be shared and intercepted. Biometric authentication may not require a physical token or memorized knowledge; the user makes a "claim of identity" and proves that claim through a biometric scan. Typically, a claim of identity takes the form of entering a username, presenting a physical badge, or presenting a passport/ID card. If the biometric scan matches a previously stored sample, the user is authenticated.

Biometrics can also distinguish an individual from a pre-defined group, which is known as identification. Identification does not require a claim of identity from the individual in question; the system's goal is to determine if the individual's identity is known to it within a specified degree of accuracy. Identification is typically conducted over a pre-defined population of users who have enrolled a sample of their biometrics.

A biometric system can only compare a current sample with a previously enrolled sample set. Before users can be authenticated using biometrics, they must enroll samples in the authentication system. It is critical to perform quality checking at the time of enrollment, as a poor quality biometric enrollment negatively affects the performance of the overall system.





Understanding Biometric Technology

A biometric is any measurable aspect of a human's physiology that can be reliably captured and used as a distinguishing identifier for that person within a defined population. Examples of biometrics used in practice are fingerprints, iris scans, hand-geometry analysis or facial recognition. A biometric security system typically uses a biometric to replace a password to authenticate an individual. However, biometrics can also be used to identify if an individual is a member of a user group (termed identification).

Like any relatively new technology, the biometrics industry is rife with hype and speculation. Common myths about biometrics include:

Myth 1: Biometrics are foolproof. No security technology is considered foolproof, especially when pitted against a determined attacker. Biometric matching of individuals is a probabilistic system, and there is a probability of a false match, where an intruder is accepted as having a valid identity. The goal is increasing the probability of a valid match and reducing the probability of a false match to as close to zero as possible.

Myth 2: Biometrics are more secure than other forms of authentication. While the biometric matching step could be considered a "secure" verification of a person's identity, a biometric-based security system is only as strong as its weakest link. Similar to attacks against cryptographic systems, attackers may not only subvert the biometric component but also an element in the supporting infrastructure, for example the storage of the biometrics.

Myth 3: A fingerprint is a secret and unique identifier. Although fingerprints are thought to be unique, an individual's fingerprints are not secrets. We leave thousands of impressions every day of our fingerprints as we move through the world.

Myth 4: Biometric systems have perfect accuracy. Biometric matching is a probabilistic algorithm; there can be no absolute certainty that two biometric samples match, merely degrees of confidence in a match. Factors such as environment, device robustness and the demographics of a population each affect the total accuracy of a biometric system.

Myth 5: Biometrics can pick terrorists out from a crowd. The technology for face detection from moving video streams is available today, but the quality of the images captured is not sufficient for accurate matching. Even under ideal lighting and camera conditions, people can easily obscure or change their face with hats, scarves, or glasses. For example, even if a face matching system had a 1% false detection error rate (an extremely optimistic scenario), in a busy metropolitan airport it would identify more than 1,000 innocent people as terrorists.

Table 3-4
Overview of strong user and device authentication solutions

Strong Authentication Solution	Description
Strong User Authentication	
Hardware Tokens	Hardware tokens are Liquid Crystal Display (LCD) panel devices that display number sequences that change periodically, for example, once per minute. In combination with a PIN, the token's software uses these sequences to create one-time passwords. Some tokens challenge the user with a built-in numeric keypad to calculate the passwords. Examples are the tokens from RSA (SecurID), ActivIdentity and Vasco.
Smart Card-Based Tokens	Smart cards are devices that can take a number of different physical forms. Most smart cards are similar to credit cards, with the addition of small, dime-sized memory chips or microprocessors. USB tokens can operate similarly to smart cards, and some vendors have implemented smart card functionality on cell phones and PDAs. Smart cards and other tokens are tamper-resistant devices that can be used for secure storage of private keys, passwords, and other personal information. Some models perform private key operations (generation, signing, and decryption) in a safe, isolated manner on the card itself. Smart card solutions require smart card readers to be deployed or integrated with the devices.
Software Tokens	Software tokens operate like hardware tokens, except that a software program installed on a user's workstation or other computing device (PDA or Pocket PC) provides a token generator or challenge/response system.
Biometric Authentication	Biometric authentication mechanisms match a physical characteristic of a user against a database record. Common methods include iris, palm, or fingerprint scans, as well as voice authentication. After years of development, these systems are becoming more reliable, yielding fewer false positives and false negatives. Prices are also falling, making biometrics increasingly practical, though still far more expensive than free passwords. Biometric solutions are particularly successful in physical facilities authentication and government applications like border security and law enforcement.
Strong Device Authentication	
Radio Frequency Identification (RFID)	An RFID system is a tag, which contains a miniscule computer chip and an antenna, that is attached to or embedded in an item. Items can be anything from a computer, to a dishwasher, to a living being. The tag transmits a signal to an electronic reader, which associates the signal with the specific item to which it is attached. The evader transmits this information to servers that collect and organize the data for tracking. RFID systems have the power to dramatically refashion such processes as the supply chain by making them more efficient, and they can bring direct consumer and societal benefits such as personalized shopping, medical reminders, and the ridentification of toxins before they reach landfills. However, the potential to tag and track every item raises privacy and civil liberty concerns. RFID technology has the potential to invade customer privacy and diminish customer control over personal information.
Trusted Platform Module (TPM)	A TPM is an embedded security chip uniquely bound to a single computer platform that can be used for both user and device authentication. TPM core components are an RSA engine, a hash engine, a key generator, and a Random Number Generator (RNG). The TPM architecture has been defined by an industry body called the Trusted Computing Group (TCG).

Why Use Biometrics?

Biometrics are appropriate for situations that require strong user authentication, in place of a password or smart card. For example, biometrics can be used as physical access controls for facilities or logical access controls for desktop computer login.

What Types of Biometrics are Available?

There are many different forms of biometrics that can verify an individual's identity. No one biometric can satisfy all operational requirements; some are more suited for office environments, and others apply to a broader range of environments.

• Fingerprint: The most recognized form of biometric is a fingerprint. Fingerprint scanners can take whole-hand images, multiple finger images, or single-finger images.

• Iris scan: Iris scanning is often confused with retinal scanning. The iris is the visible colored part of the eye that surrounds the pupil. It is believed to be unique for each human.

• Retinal scan: Often confused with iris scanning, the retina is an area at the rear of the eyeball. An iris scanner is a specialized device that must scan the eye from very close range. Despite its accuracy, retinal scanning is rarely used because of privacy and usability issues.

• Vein pattern: The pattern of veins in the hand is believed to be unique across the population. The hand is placed on a reader and illuminated with infrared light to detect the veins.

• Hand geometry: The geometry of the hand is potentially unique across the population. The hand is placed on a scanner and various measurements are taken of finger spacing, length, and angle.

• Voice analysis: Voice biometrics analyze the inflections of a speaker's voice to authenticate the speaker.

• Face: Facial biometrics analyze a picture (or stream of images) from a human face. Typically, a facial biometric system uses a standard digital camera to take a picture of a face.

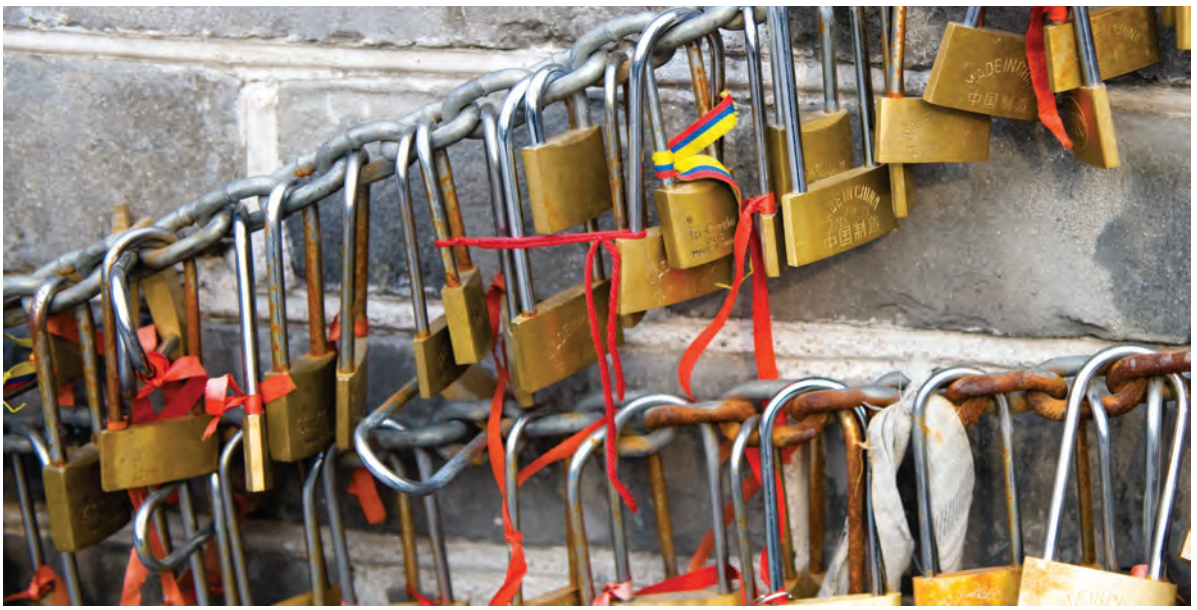

Biometrics Issues
Biometric technology brings a different set of operational and security issues into consideration:

• Ease of measurement: A fingerprint is easy to measure in an office environment, but it may be difficult to capture in an industrial setting where operators wear protective gloves. If the biometric is not easy to measure, users will be frustrated.
• Range of accuracy: Biometrics accuracy is highly dependent on the choice of sensors, the user population, the operational environment, and the manner of use.
• Device robustness: Biometric readers are subject to the usual wear and tear any office product experiences. Some will fail more rapidly than others, and each device may fail in different ways. For example, some readers may degrade the image that they capture so that a comparison may still be made, but others may simply refuse to capture an image.
• Technology improvements: The biometrics industry is moving at such a rapid pace of change that choosing a device and biometric system is often a moving target. When investing in biometric systems today, organizations should leave options open to move to a newer technology in the future.
• Acceptability to users: If users feel that the biometric system is slowing them down in their business activities, they will attempt to find ways around the system, or will claim that it does not work correctly.

HP Strong Authentication Solutions
HP offers a variety of strong authentication solutions. These include smart card technologies, biometric devices, TPM solutions and RFID technologies. Smart Card Security for HP ProtectTools is HP's strong authentication solution rooted in smart cards. It has several unique features:

• Smart card in BIOS allows for pre-boot authentication and is OS independent.
• Smart card logon allows for strong, smart card-based Microsoft Windows authentication without requiring a PKI.

Smart Card Security for HP ProtectTools is available on select commercial desktop, notebook, and workstation models. For the current model list, see **www.hp.com/go/notebooks** or **www.hp.com/go/desktops**.

HP biometric solutions include the HP USB biometric fingerprint reader, the HP built-in fingerprint reader on select models of the HP iPAQ Pocket PC, and the integrated biometric readers on select commercial notebooks. Credential Manager for HP ProtectTools, a module of HP ProtectTools, is standard on select commercial notebooks, desktops and workstations. Credential Manager offers an efficient and integrated means of managing multi-factor authentication with a fingerprint reader, smart card or token.

Tying the TPM to authentication technologies provides even stronger protection of identities by connecting software to hardware. Embedded Security for HP ProtectTools uses the TPM embedded security chip to help protect against unauthorized access to sensitive user data and credentials. It is supported on all HP business notebooks, and select desktops and workstations configured with a TPM embedded security chip. See **www.hp.com/go/notebooks** or **www.hp.com/go/desktops** for specific information on platforms. HP is the co-founder and leader of TPM specification development within the TCG. For more information about the TCG, see **www.trustedcomputinggroup.org**.

HP is leveraging its experience, expertise, and understanding of RFID technologies in its consulting services, and HP is one of the world's largest players in the space. HP has taken a leadership role in developing RFID standards. In fact, HP was performing trials and pilot projects long before January 2005, when Wal-Mart and the U.S. Department of Defense insisted that their suppliers begin using RFID tagging. With the world's ninth largest non-military supply chain, HP wants to demonstrate the value of RFID systems for increasing the end-to-end speed and visibility of supply chains.

SSO is the ability for a user to authenticate once to a single authentication authority, obtain a credential token or artifact with a defined lifespan, and use it to access other protected resources without re-authenticating.

8.2.1.2. Single Sign-On (SSO)

Single sign-on is the ability for a user to authenticate once to a single authentication authority, obtain a credential token or artifact with a defined lifespan, and use it to access other protected resources without re-authenticating. The Open Group (**www.opengroup.org**) defines SSO as the "mechanism whereby a single action of user authentication and authorization can permit a user to access all computers and systems where that user has access permission, without the need to enter multiple passwords."

HP and SSO

For enterprise and web SSO, HP takes a best-of-breed approach, leveraging the solutions from industry leaders such as Citrix Systems (Password Manager; **www.citrix.com**).

For facilitating SSO on the client side, HP offers Credential Manager for HP ProtectTools. Credential Manager for HP ProtectTools is a client-side credential caching-based SSO solution. Credential Manager acts as a personal password vault that makes accessing protected information more secure and convenient by automatically remembering credentials for websites, applications and protected network resources. Thanks to Credential Manager for HP ProtectTools, users no longer need to remember multiple passwords. Additionally, it provides enhanced protection against unauthorized access to a commercial notebook, desktop, or workstation, including alternatives to passwords when logging on to Microsoft Windows.

For identity control and access protection, Credential Manager for HP ProtectTools allows the user to define and implement multi-factor authentication capabilities. For example, when authenticating to a PC, users can be required to provide combinations of different security technologies, such as a smart card or a biometric ID. Furthermore, Credential Manager for HP ProtectTools password store is protected via encryption. It can also be hardened through the use of a TPM embedded security chip.

8.2.1.3. Authentication Support in HP-UX 11i

The HP-UX 11i OS is designed to meet the security requirements of demanding environments. With a pluggable framework for authentication, HP-UX 11i can integrate into security infrastructures and also maintain a pervasive management solution. The HP-UX Pluggable Authentication Module (PAM) subsystem provides a pluggable authentication backbone for secured authentication services on HP-UX. HP offers several authentication pluggable security modules for PAM, including integration with Kerberos and LDAP.

The HP-UX 11i AAA Radius server can act as the front end to the identity management system by operating at the point of entry to a network (the access control point). When the HP-UX 11i AAA Radius server is tied to the Red Hat/Netscape Directory Server, external remote access is authenticated. In addition, this arrangement controls and passes access to network usage accounting systems and eventual billing software. This configuration is especially useful for Telco and Internet service providers.

8.2.2. Authorization Technologies

The goal of an authorization system is to protect resources and information while allowing fluid access for legitimate users of these resources. Authorization is the act of granting subjects access rights to protected resources. The main difficulty is scaling authorization policy administration to thousands - possibly millions - of subjects and protected resources. As the numbers grow, administrators need to reduce the ratio of policies to the number of subjects and protected resources without compromising the security of the system.

The main difficulty is scaling authorization policy administration to thousands - possibly millions - of subjects and protected resources. As the numbers grow, administrators need to reduce the ratio of policies to the number of subjects and protected resources without compromising the security of the system.

Authorization policies are rules for determining which subjects are allowed to access resources. In some cases, privacy considerations may require support for some form of anonymous or pseudonymous access. In most cases, however, users must be identified prior to receiving the authorization to access resources. An identity management infrastructure is therefore critical to establishing users' identities as the basis for authorizing access to resources.

Two interesting access control models used in the identity management infrastructure are the role-based access control (RBAC) model and the rule set-based access control (RSBAC) model. A role is an organizational job function with a clear definition of inherent responsibility and authority (permissions). The process by which an enterprise develops, applies, and maintains RBAC is known as role engineering. As old roles are retired or modified and new roles are defined to meet changing business needs, an enterprise defines processes for updating roles.

Role-based approaches are suitable when job functions are easily partitioned. Wide-scale implementations remain stalled because of the complex nature and large scope of role engineering projects, transitory job assignments in knowledge-based organizations, lack of funding, limited standardization, and proprietary access control mechanisms. A common challenge facing role-based systems is finding agreement among stakeholders for standardized vocabularies and role definitions.

An alternative to an RBAC approach is RSBAC. With this approach, access is granted or denied based on a set of pre-defined rules or organizational policies. Access control decisions can change dynamically based on access control policies. Rules are context-dependent: they can take into account things like the time of day, resource type, and access location.

8.2.2.2. Authorization Support in HP-UX 11i

HP-UX 11i is designed to meet the security requirements of demanding environments and provides a pluggable framework for both authentication and authorization. As part of HP-UX 11i v2, HP-UX adds the Access Control Policy Switch, which is available as part of the RBAC subsystem.

With the PAM, RBAC, and identity management integration features of HP-UX 11i, administration of authentication and OS- or application-specific privileges can be centrally managed through identity management products. These features enable the pervasive management capabilities required by today's organizations.

As noted previously, the HP-UX 11i AAA Radius server can act as the front end to the identity management system by operating at the access control point. When combined with the Red Hat/Netscape Directory Server, the system offers additional benefits for authentication, access control, and accounting that are especially useful for telco and internet service providers.

Increasingly, regulatory demands require enterprises to understand what their users are doing. The challenge is consolidating and making sense of identity data with respect to policies and regulations in a complex and ever-changing environment.

8.2.3. Auditing Technologies

Increasingly, regulatory demands require enterprises to understand what their users are doing. The challenge is consolidating and making sense of identity data with respect to policies and regulations in a complex and ever-changing environment. Auditing systems capture security-related events in identity management systems and ensure accountability for the underlying IT and security infrastructure. Complete and accurate audit and event records provide the evidence that enterprises need to demonstrate compliance with business, security, legal, and regulatory mandates. It is especially critical to audit the authorization, provisioning, and privacy components of identity management systems, which may create or remove user privileges and accounts.

Typically, it is difficult and costly for an organization to determine who did what and when. There are not easy methods for management to determine compliance with specific regulations or to view trends. Additionally, organizations cannot track who knew about violations - clear and specific ownership does not exist.

As a result, organizations risk failing regulatory audits. This is a critical concern for all CIOs, CSOs, and personnel directly responsible for the protection of organizational resources, including privacy officers, risk managers, and auditors. Company executives can now be personally liable for failures in their control systems, and regulatory failures can lead to critical business impacts such as lost customers, fines, or jail time for those responsible.

Challenges requiring organizations to better manage audit and reporting include the need to:

- Comply with regulations such as the Basel II Accord, EU Directive on Data Protection, Japanese Data Protection, and the Personal Information Protection and Electronic Documents Act (PIPEDA)
- Control the identity infrastructure; specifically, knowing who has access to what and who approved access
- Track and trend compliance for repeatable success with audits, for example, compare compliance for today, last week, the last six months, and longer

HP uses a best-of-breed and customer preference approach for recommending multi- platform and multi- application auditing solutions to customers.

8.3. Privacy Management

In the context of electronic privacy, users express concerns regarding today's IT systems and environments like the Internet. Some of the key privacy concerns include:

- Data is collected silently. This is facilitated by the Web, which allows large quantities of data to be collected inexpensively and unobtrusively.
- Data from multiple sources may be merged. Non-identifiable information can become identifiable when merged with other sources and information.
- Data collected for business purposes may be used in civil and criminal proceedings.
- Data collected for business purposes may be forwarded to third parties, without notifying the user.
- Data may be copied and used without authorization. This may happen simply when data are used for purposes other than those for which they were collected; it may also happen as a result of failure of access controls allowing malicious parties, within or outside the collecting organization, to obtain data for criminal purposes such as fraud and identity theft.
- Users are not given meaningful choices for the use of their personal data.




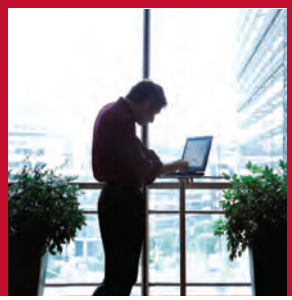


Trends Topic
Global Privacy

Privacy involves the treatment of personal data. A key concern of individuals who provide personal data is its authorized but undisclosed secondary use. Preventing undisclosed use requires the collecting party to inform individuals of any intent to share their personal data beyond the primary use. The collecting party should allow individuals a choice to opt in and then subsequently respect the agreement. These circumstances bring the matter of privacy closer to the contractual and legal spheres rather than the technological.

Legal and cultural approaches to privacy differ widely around the globe. The European Union, for example, has explicit data-protection regulations. In the U.S., individual privacy is generally a matter of contractual and informed-consent practices, although regulations apply in specific sectors such as the health industry. Privacy approaches in Asian countries vary widely in both legal requirements and prevailing attitudes. Further complicating the picture are tensions related to anonymity, such as preserving the right to individual privacy while preventing the use of anonymity as a means for criminals to conceal activities, and allowing law enforcement and homeland security agencies to gather extensive personal data to pursue preventive and investigative programs.

Some specific technologies that address global privacy and anonymity pressures are in development. For example, there are many cryptographic techniques for "managed anonymity" or pseudonymity. These techniques, such as double-spending detection in electronic cash schemes, protect individual identities for normal operations; however, they can reveal identities in exceptional circumstances. Privacy-preserving computations are also feasible in some cases. For example, the holders of two different databases can jointly discover common entries without revealing information to each other. The provision of reliable remote execution environments, where executable content can run on a machine that demonstrates its trustworthiness, offers the promise of data transmission with continued effective controls on data use.

Privacy considerations typically arise when an organization needs to collect and store customer, employee, and other private data. Other situations that raise privacy concerns include demonstrating conformance with privacy regulations and forwarding user information (identity information, web-service access information, security assertions, or localization data) to third-party service providers.

Numerous efforts have produced legislative frameworks for privacy. Examples are the EU Directive on Data Protection, U.S. laws such as HIPAA and the Children's Online Privacy Protection Act (COPPA), and frameworks such as Safe Harbor. However, privacy and data protection laws are hard to enforce when personal information spreads across boundaries. In general, users have little understanding or knowledge of privacy laws and their implications.

Privacy management in an IT environment has many different aspects. These include negotiation, policy lifecycle management, enforcement, monitoring, decision support, violation detection, preserving computation, data minimization and transformation, rating and branding, verification and certification, auditing and accountability, mediation and delegation, anonymization and pseudonymization, and user training. In the context of identity management solutions, privacy-protecting technologies can be viewed as an extension to 1) authorization systems and 2) provisioning and identity lifecycle management solutions. Authorization policies control data access based on factors like privacy regulations and user consent. Obligation policies control how to handle the lifecycle of identity information during its lifetime, including data deletion/retention, data transformation/minimization, notifications, etc.

Privacy management is an identity management area that requires much work and effort. HP and other major IT players like IBM Corporation are leading key developments in the privacy management space. HP is involved in several cross-industry and government-driven privacy standardization initiatives. HP Labs, jointly with HP business groups (and also in the context of international projects, such as the European Union (EU)-funded PRIME project), are actively researching and developing technologies and solutions in the privacy space. HP Labs privacy research has been focusing on a Privacy Policy Enforcement System and an Obligation Management System to model, deploy and enforce privacy policies and obligations at the operational level: the feasibility of these R&D systems has been demonstrated by showing how they can be integrated with HP identity management solutions.

HP Labs privacy research has been focusing on a Privacy Policy Enforcement System and an Obligation Management System to model, deploy and enforce privacy policies and obligations at the operational level.

HP Labs Privacy Policy Enforcement System

Traditional access control solutions (that involve users, their roles, protected resources and access rights) are necessary but not sufficient to enforce privacy constraints when accessing personal data. These solutions need to be extended to keep into account the purpose by which data has been collected, consent given by data subjects and other conditions.

This HPL research focuses on the development of a privacy-aware access control system that enforces privacy policies (defined by privacy administrators and based on data subjects' privacy preferences) on personal data stored in heterogeneous enterprise data repositories. In this system, privacy policies explicitly define the purposes for which personal data can be accessed, how to keep into account data subjects' consent and what actions need to be fulfilled when the data is accessed (filtering out data, blocking access, logging, etc.).

The HP Labs Privacy Policy Enforcement System provides the following key functionalities: (1) it allows administrators to graphically author policies involving both privacy and access control aspects; (2) it allows for fine–grained modeling of personal data (stored in relational databases, LDAP directories, etc.) that are subjected to privacy policies; (3) it allows for the deployment of policies and the decision–making process based on these policies; (4) it allows for the enforcement of privacy–related policies when the data is accessed.

HP Labs Privacy Obligation Management System

Access control solutions cannot deal with all aspects of privacy policy enforcement. In particular they are not designed to handle constraints dictated by privacy obligations, such as duties and expectations on data deletion, data retention, data transformation, etc. For example, data might need to be deleted after a predefined period of time, independently from the fact that this data has ever been accessed. Privacy obligations introduce the need to deal with privacy–aware information lifecycle management - i.e. ensure that the creation, storage, modification and deletion of data is driven by privacy criteria.

HP Labs has been working on this area, in particular on the problem of explicitly representing obligations, reasoning on them, enforcing and monitoring them. A main differentiation of the HP Labs work is the clear separation between access control management and obligation management – without imposing a subordinated view of obligations to access control policies.

HP Labs has defined a privacy obligation management modeland an Obligation Management System (OMS) to explicitly manage privacy obligations on personal data.

The HP Labs OMS provides the following functionalities: (1) explicit representation of privacy obligations as reaction rules; (2) scheduling of privacy obligations; (3) enforcement of obligations; (4) monitoring of enforced obligations. Privacy obligations can be automatically derived from privacy preferences (e.g. requests for deletion or notifications) expressed by people or administrators on personal data. These obligations are scheduled by the OMS system based on relevant events. If triggered by these events, OMS enforces privacy obligations, for example by deleting data, sending notifications or triggering workflows. Enforced obligations are monitored for a predefined period of time for compliance reasons.

A full working prototype of the HP Labs OMS has been implemented in the context of the EU PRIME Project (see **www.prime-project.eu**).

8.4. Identity Lifecycle Management

Identity lifecycle management or provisioning solutions are similar to SSO solutions in that they operate from the top down; the application manages all of the systems under it. Administrative functions - from the essential add, modify, and delete to the more general maintenance and monitoring - are under the control of the provisioning system. Provisioning functions can also include non – electronic tasks such as identifying a cubicle, connecting a network port, and acquiring a PC.

With provisioning products, organizations risk implementing one solution that can potentially clash with another. Provisioning solutions often incorporate other parts of the identity management framework, such as self – service and password management.

The only mature provisioning – specific standard at this time is the Service Provisioning Markup Language (SPML). SPML messages facilitate the creation, modification, activation, suspension, enablement, and deletion of identity – related data in different identity repositories. The Organization for the Advancement of Structured Information Standards (OASIS) has been working on the SPML specification since late 2001. For more information, see **www.oasis–open.org/home/index.php**.

HP uses a best-of-breed and customer preference approach for recommending identity lifecycle management and provisioning solutions to customers.

8.5. Federated Identity Management

Federation enables the trusted interchange of security-related information between different autonomous policy domains. Security-related information includes authentication, authorization, and auditing data. Although federation is generally used in the context of an inter-enterprise security mechanism, it can also be used within an enterprise to provide tighter integration between loosely coupled ecosystems.

Typically, a federation provides a common framework for trust - a standard syntax, vocabulary, attribute set, and set of policies and practices for the trusted interchange of security-related information. Bilateral (federation) agreements between partners are often required to negotiate the specifics of access, such as which users or systems can access which resources, under what circumstances, and under what contractual relationships. Access control always remains with the owner of the resource. A federation might also define minimum acceptable trust levels or authentication mechanisms required for specific circumstances.

A federation agreement always deals with two entities: an asserting party that generates security assertions and a relying party that trusts the security assertion made by the asserting party. There are a number of federations being formed, supporting a variety of vertical marketplaces, communities of interest (financial services, health sciences, research and education), and geopolitical boundaries (state and national governments).

HP uses a best-of-breed and customer preference approach for recommending identity federation solutions to customers.

Figure 3-5
How the identity federation standards stack up

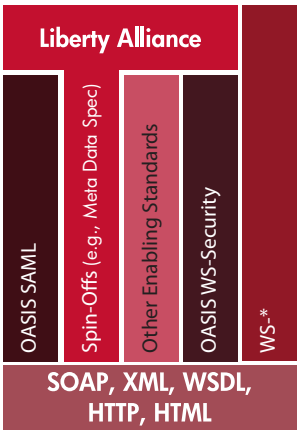

8.5.1. Federation Standards

A variety of standards, specifications, and protocols relate to federated identity management. Figure 3-5 shows the positioning of some of the relevant federated identity management standards. The Liberty Alliance specifications define the protocol messages, profiles, and processing rules for identity federation and management. They rely heavily on other standards such as SAML and WS-Security. Additionally, the Liberty Alliance has contributed portions of its specification to the technical committee working on SAML. More information is available from **www.projectliberty.org**. HP endorses the Liberty Alliance and actively participates in the creation of its specifications.

SAML is an OASIS specification that provides a set of XML and Simple Object Application Protocol (SOAP)-based services, protocols, and formats for exchanging authentication and authorization information. More information is available from **www.oasisopen.org/committees/tc_home.php?wg_abbrev=security**.

WS-Security is another OASIS specification that defines mechanisms implemented in SOAP headers. These mechanisms are designed to enhance SOAP messaging by providing a quality of protection through message integrity, message confidentiality, and single message authentication. More information is available from **www.oasis-open.org/committees/tc_home.php?wg_abbrev=wss**.

The Web Services protocol specifications (WS-*) are currently in development by Microsoft Corporation and IBM Corporation. They include specifications for WS-Policy, WS-Trust, and WS-Federation.

Other identity management enabling standards are:
- Service Provisioning Markup Language (SPML), **www.oasisopen.org/committees/tc_home.php?wg_abbrev=provision**.
- XML Access Control Markup Language (XACML), **www.oasisopen.org/committees/tc_home.php?wg_abbrev=xacml**.
- XML Signature, **www.w3.org/Signature**.
- XML Encryption, **www.w3.org/Encryption**.

8.6. HP's National Identity System

The HP National Identity System (HP NIS) is a public-sector identity credentialing solution that provides the component modules required to implement a national-scale enrollment and document issuance solution. HP NIS supports workflow processes for citizen enrollment, establishes and maintains a distributed citizen registry system, manages identification document processes, and incorporates biometric and PKI verification.

8.6.1. Multiple Stakeholders

There are multiple participants in a national ID system, and each group brings its own requirements and perspective. A successful credential issuance system must satisfy the needs of each stakeholder. The three stakeholder classes for HP NIS are:

- Customers: National governments or government departments are the purchasing customers for most national ID or passport systems.
- Suppliers: A single vendor cannot offer a comprehensive national ID solution; for example, there are specialists in passport production, ID card printing, and smart cards that must work together to meet customer needs.
- Citizens: The ultimate credential holders and users are citizens who expect their government to provide a secure document that is durable, easy-to-use, and inexpensive.

Figure 3-6
The identity document lifecycle

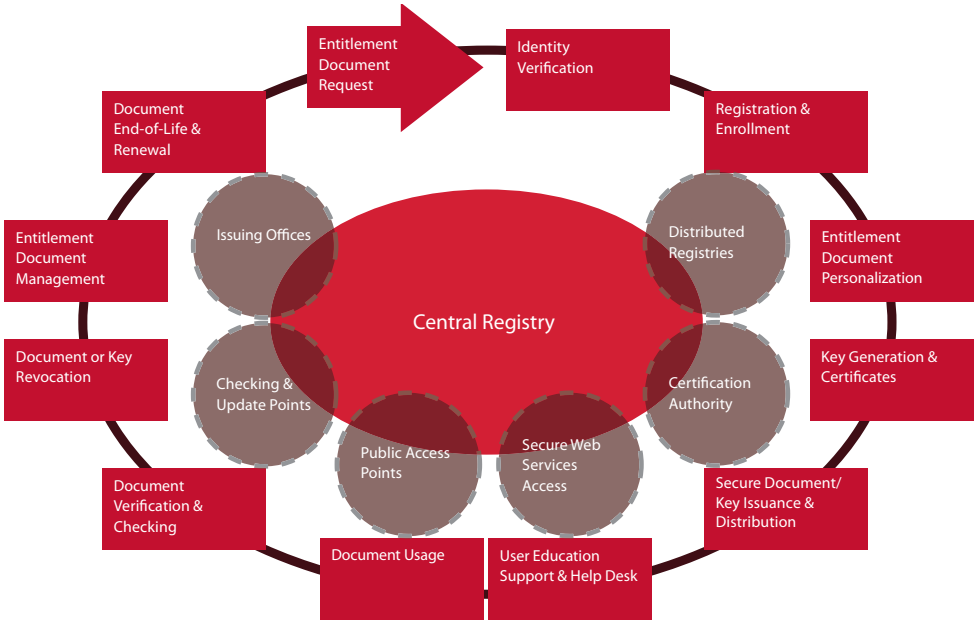

8.6.2. The identity document lifecycle

Identity documents follow a well-defined lifecycle from issuance through end-of-life disposal, as illustrated in Figures 3-6 and 3-7. The lifecycle mirrors the processes in place in the government department that issues the document. For example, the lifecycle for a passport document includes document request (applying for a passport), adjudication (determining if a passport should be issued), issuance (printing the passport and encoding data on the chip), and usage (using the passport to cross the border).

What is a public-sector identity credential?

A public-sector identity credential is a document issued by a government or state agency used by a citizen or resident to gain access to state provided services. Examples of these identity credentials include ID cards, health/welfare entitlement cards, medical information cards, passports, visas and driver's licenses.

An identity credential issued by a government is often considered a secure document and thus requires security features to reduce the likelihood of fraudulent copies being produced, or existing documents being modified. Furthermore there must be trust in the integrity of the document issuance process so that suitable background checks are performed before a document is issued. For example, in countries which issue ID cards to citizens there is an expectation that one card is issued per citizen and each citizen can hold only one ID card. Machine-readable travel documents include passports (both traditional and ePassports) and must comply with international standards for data encoding and formatting to permit inspection at border crossing points.

Figure 3-7
Central registry

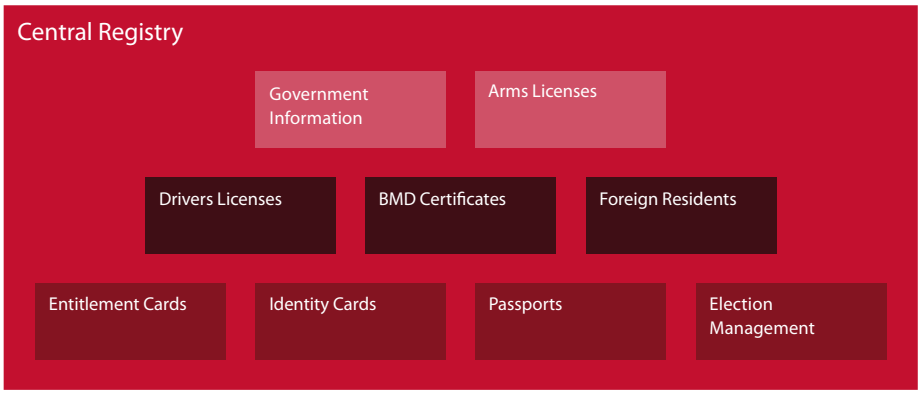

8.6.3. HP NIS Architecture

NIS incorporates several high-level core processes consisting of enrollment, validation, registration, issuace, card usage, support, and lifecycle management. The processes begin after properly validating an applicant in accordance with customer policy.

Enrollment

Enrollment typically includes review and authentication of an applicant's official breeder documents, personal information, and biometric data. Customer policy dictates the number and types of breeder documents to authenticate and the nature of the authentication process. Generally, an apparatus electronically screens these documents to detect document tampering (such as an altered driver's license with a different photograph, signature, or other demographic information). The analysis of several factors easily detects tampering, for example, changes in laminating materials, changes in text density and color, or minute misalignments of background patterns or holograms. Checking documents such as birth certificates with the issuing agency for authenticity is another common requirement.

Demographic and biometric information may be collected from applicants. Demographic information may be uploaded via optical character recognition technology from paper forms completed by applicants; entered by applicants at kiosks or workstations located at application processing centers; or entered, in part, from a web-based interface. Biometric scanning occurs at an application processing station. All information, demographic (from whatever input source) and biometric, is encrypted and transmitted to a secure enrollment database where it remains in encrypted format. Once the applicant authenticates the demographic information gathered during the enrollment process, the application document is signed. Signed documents are encrypted and forwarded to the enrollment database.

All information submitted by the applicant is stored exactly as submitted (either text input or scanned document). Subsequent changes are captured and stored as submitted, and the system maintains a complete history of all data entries. Data management of applications is configurable according to local legislative requirements. Receipt of the signed application form is the trigger for the next step of validation.

Validation

Validation is an automated process that electronically verifies demographic and biometric information with existing watch lists and with public records on file in various governmental and targeted public databases. Biographic information, such as fingerprints, are checked against existing national registries and against the NIS database of completed registrants to ensure that an applicant is not registering with an illegal alias to obtain a fraudulent credential. Applicable demographic information, such as birth date, address, and any other information designated by customer policy, is crosschecked against external databases and registers. During processing, the applications database maintains a historical record of validation processing.

When validation is complete, the applicant's file is considered a positive verification, or it is listed as a suspicious record for further examination and adjudication. Customer policy dictates the review process; generally, it requires human intervention. If desired by the customer, however, the validation review can be entirely automated.

Registration
Validated records are forwarded to the personnel registry (database), where a complete record of all data files, including biometrics, are stored and maintained. Separately, a dedicated biometric database stores biometric data files to provide visual and electronic replication. Records from this database help create digital information for the identification document. The biometric records also authenticate online requests for identity verification when cardholders attempt to gain physical or logical entry into locations and/or systems.

Issuance
Issuance is an automated four-step procedure:

1. Personalization: In this first step the blank document is printed and digital information is loaded if a chip is present on the card. Examples of information that may be stored on a chip include applets, certificates, key materials, demographic information, and biometric information for visual or electronic viewing.
2. Quality assurance: The finished document is visually and electronically checked to ensure that it complies with requirements. The visual inspection process is usually automated.
3. Issuance: The document is issued to the applicant, often requiring an identity check.
4. Activation: If the card contains a chip, it must be activated by the holder before it can be used.

The exact format, appearance, and electrical characteristics conform to the standards directed by customer policy.

Personalization
During personalization, information together with security features are physically printed on the card. The card's chip is electronically formatted with personalization data, PKI materials if required, certificates if required, and other customer-specified codes such as PINs and/or personal unlock keys (PUKs). Personalization artifacts physically printed on the card typically consist of the card holder's name, picture, signature, and card expiration date. The same information is electronically stored on the card's chip, together with biometric information such as face or fingerprints images.

Quality Assurance
Credential quality assurance is an automated procedure to verify that:

• The data printed on the card conforms to the precision and security requirements specified by the customer (e.g. all elements are clearly printed).
• The data encoded on the chip is correctly encoded.

Issuance
Detailed procedures for the handover of an identity card to a user conform to customer policy. HP NIS is a versatile system designed to support a multitude of options including the use of mailing/courier services, internal document distribution systems, and in-person card handover.

Activation
Activation of the identity card can occur upon handover (if it is performed at a site with a card reader) or upon first use at a site with a card reader. Activation of the card and all subsequent usage is recorded and maintained in an auditable registry per the customer's policies and procedures.

Card Usage
During card usage, a comparison process determines that the credential holder is the person depicted by the information on the credential document. Depending upon the level of security employed, the comparison process can include:

• A visual comparison of the credential holder's face with the photograph on the credential
• An automated comparison of live biometric data (for example, from a fingerprint scanner) to biometric data stored on the credential document
• A comparison of live biometric data with a remote biometric data registry, and an authentication of the PKI materials by a central database

Support
Credential holder support includes both a self-administration component and a help desk component. The self-administration component lets end users make limited credential changes, such as modifying PINs, and changes to non-critical data (as defined by customer policies). Help desk staff can block credential use in cases of lost or stolen identity cards and unblock credentials under specific situations.

Lifecycle Management

Lifecycle management is the most critical element of HP NIS. Lifecycle management is the all-encompassing maintenance and management of every piece of identity management information and material, including, but not limited to, blank identity documents received into inventory, collected demographic and biometric data, accumulated vetting reports and other ancillary information from official sources, active identity credential status, revoked or suspended credentials, and auditing reports.

HP NIS lifecycle management consists of six general processes:

- Track and manage all unused security consumables, such as blank identity cards, PKI materials, certificates, holograms, and laminates. The inventory and control function tracks quantities and locations of these materials, and it generates reorder notices as appropriate.
- Track and manage identity card OS and applet lifecycles, including revision and version control for all installed firmware and software, by credential serial number.
- Manage application forms for all materials submitted to NIS during the enrollment of an applicant. Submitted materials are never deleted.
- Manage the personnel registry and biometric database where copies of all validated documents, reports, and other data entries are stored, together with biometric data files, respectively. Once entered, items are not deleted from the personnel registry.
- Track and manage all issued physical credentials, including a complete audit trail for credential enabling, use, suspension, query, revocation, reissue, loss, recovery, and compromise of the issued document.
- Track and manage all personalization entries, consisting of all initial and updated demographic, biometric, PKI materials, and/or certificate entries on the issued credential document.

8.6.4. NIS Governance

Governance refers to implementing and enforcing the policies that control data encryption, data storage, and access to stored data. NIS is designed to meet or exceed the stringent security standards dictated by governments around the world. All sensitive data is encrypted and signed before transmission and when stored on any permanent or erasable medium. Encryption applies to all data, including data collected and stored during enrollment and data permanently stored in the personal and biometric registries.

8.6.5. Confidence in Identity Documents

For an identity document to be useful, both the issuer and the validator of the document must have confidence in its integrity and authenticity. Document integrity applies to the physical document and any electronic data encoded on the document's smart card chip. Physical security features include security foils, laminates, microprinting, and holograms. These features allow a visual inspection of the document to detect tampering or forgery. Logical security features include digital signatures on the chip contents, encrypted sessions between the chip and a reader device, and PKI technology to ensure that the electronic data on the chip was issued by a known entity (for example, a country or state).

At the least secure level, an individual seeking entry into a restricted space surrenders his or her identification card to an attendant, who verifies that the card is not altered, that the photograph on the card is that of the presenter, and that the card has not expired. Once satisfied that identity is verified, the attendant grants access to a controlled area. This example checks only the physical security features of the identity document.

Security is increased when individuals use a card reader for access. In this example, individuals pass their card through a reader and enter a PIN to verify they are the cardholder. The card reader verifies that the PIN matches the stored PIN on the card and that the card has not expired. If the data matches, individuals gain physical entry into the controlled facility or electrical access to a computer or network. To further increase security, an attendant can also check the physical security features of the identity document.

Incorporating biometric authentication with the card reader substantially increases security. Individuals seeking entry into a controlled area or access to a computer or network must pass their card through a reader and enter a PIN for comparison against the stored PIN. Additionally, individuals must submit a biometric sample (such as a fingerprint scan) to compare with the digital biometric file stored on the card. Finally, the reader verifies that the card is not expired. If all checks are satisfactory, the individual is granted physical entry or logical access. Again, the process is more secure if an attendant verifies the picture on the ID card and checks that the card is not altered.

8.6.6. NIS Summary

HP NIS is a versatile credentialing system, designed and engineered from the ground up to meet the most demanding requirements for stringent security, worldwide. It is easily deployed as a modular system, and it efficiently operates from widely dispersed, remote locations. The system's scalability fits deployment to small, large, or global populations with equal ease. Deployment causes minimal impact on day-to-day activities and does not require excessive overhead. The ease of scalability means that NIS is capable of processing very large citizenries, including national populations, in a relatively short time span.

9. Successfully Approaching Identity Management

Identity management employs a consolidated view of an "identity" across the enterprise, including identity information and attributes aspects, authentication aspects and privilege and entitlement aspects affecting people, processes and technology.
Because of the scope and enterprise-wide reach of identity management, starting in the human resources organization, over to the business, internal and external audit and compliance and security organizations, to IT and helpdesk, identity management programs need to be set up and justified correctly in order to become successful. Identity management will then become an integral strategy of the enterprise iterating through defined and planned cycles of improvements.

9.1. Review and Envision Phase

A critical success factor of identity management initiatives is to move early toward a value consensus shared across all organizations that will be impacted in your identity management initiative. Shared value consensus will be a key to motivating separate interests to work together toward a common cause.
We help enterprises determining the readiness to launch improvement initiatives. Using the Business Readiness service for identity management, HP Services supports enterprises define targeted business capabilities and postulate impact of those capabilities to the organization.

Utilizing the experience from identity management programs in other organizations, business impact is seen key to the program and can be categorized into the following areas:

Regulation conformance: Enterprises that need to conform to regulations such Sarbanes-Oxley, Basel II or other industry specific regulations are often challenged with optimizing the regular auditing process for compliancy to required security controls.

Security: Enterprise security risk and business impact analysis results in risk mitigation strategies often in the area of inadequate account and permission management with lifecycle events of job and role changes and terminations.

Data quality: Due to the fast growth of applications and identity systems and repositories, companies have been struggling in setting up the right lifecycle process resulting in inconsistencies in profile and entitlement data.

Agility and productivity and user convenience: Identity lifecycle processes such as user onboarding, role change and termination can be complex processes with many people and organizations involved. Enterprises with high dynamics with the work force or the customer base look at automation to respond to changes quicker.

Cost Reduction: Enterprises looking at return of investment calculations reviewing spending in many areas such as the helpdesk performing manual provisioning and password management, user management efforts within business applications or the consolidation of point identity management solutions.

9.2. Definition Phase

Developing an overarching strategy represents a key success factor for the identity management initiative. HP Services provides the Discovery and Framework service to help enterprises organize this initiative, achieving a common understanding of the current situation and providing the strategic alignment throughout the organization of the desired state, including a roadmap of how and when to get there. Utilizing individual workshops with various parts of the organization supports the discovery of the current state from all angels including major painpoints existing throughout the organization.

To achieve enterprise value, the ultimate scope must deliver benefits meaningful across the constituencies, however, a broad scope can create what may appear initially to be an insurmountable project challenge. To overcome this, the prioritization and evaluation of major pain points including the alignment of corporate and IT architecture strategy is the enablement for the development of not only a high-level architecture but also the partitioning and sequence of the ultimate scope into an implementation roadmap of manageable projects in which progress can be measure and goals can be adjusted as necessary.

9.3. Design and Implementation Phase

Once the scope and project has been clearly defined for the particular cycle of improvement, the detailed functional, technical and implementation requirements will be discovered and aligned to the originating business requirements.

The resulting design must comprehend not only tools and technologies but also lifecycle processes and organizational responsibilities and agreements.

The implementation must factor in the best practices for the technology, process and organizational impact. HP Services utilizes its 10-step implementation methodology for identity management deployments to ensure successful deployments.

9.4. Identity Management Success Factors

In addition to the many success factors that are specific to the program phases, HP has developed overall key success factors for enterprises' identity management initiatives:

Manageable Project Phases: Because of the global reach of identity management, the amount of functional and technical requirements an identity management initiative is challenged with can be large. Therefore identity management programs must ensure that business value is returned to the organization in a timely manner, while not trying to attempt to solve too many issues at the same time. This includes the clear definition of:

- Scope
- Realistic timelines
- Demonstrable benefits

Governance: Operationalizing identity management programs require the establishment of strategy alignment and definition of responsibilities throughout the organization. This includes:

- Steering committee to prioritize efforts and resolve conflicts
- Accountability from all program participants
- Continuous communication and loopback to business stakeholders

Experience: Identity management programs require extensive experience in a lot of very different areas, covering:

- Business processes
- Legal and regulatory implications
- Technology
- Operations

"Building a flexible and responsive identity management foundation requires an investment of time, money, energy, and resources, but it can help enterprises gain an advantage over their competition."
-The Burton Group, "Building the Business Case for identity management Investment", v2, January 3, 2006

10. HP Identity Management Services

Based on the principles that were outlined in the section on "Successfully Approaching Identity Management" above, HP Services provides the following end-to-end service offerings for identity management (the offerings are summarized in Figure 3-8):

Business Readiness Workshop

This service analyzes and assesses the true business value that an identity management solution could have for your organization. By taking a high-level view of your current business objectives, internal processes and IT infrastructure, you'll be able to pinpoint where identity management might have the greatest financial impact for your organization; some of the key areas to target for improvement; and the most effective long-term strategy. Focal areas include regulation conformance, security, data quality, agility and productivity and cost reduction.

Identity Management Discovery and Framework Service

This service follows a proven methodology to lay the foundation for effective change of your identity management processes and IT environment. Working with you, we gather information, establish guiding design principles, help express business value, outline the identity management vision architecture and develop a roadmap for improvement phases.

Figure 3-8
HP identity management services

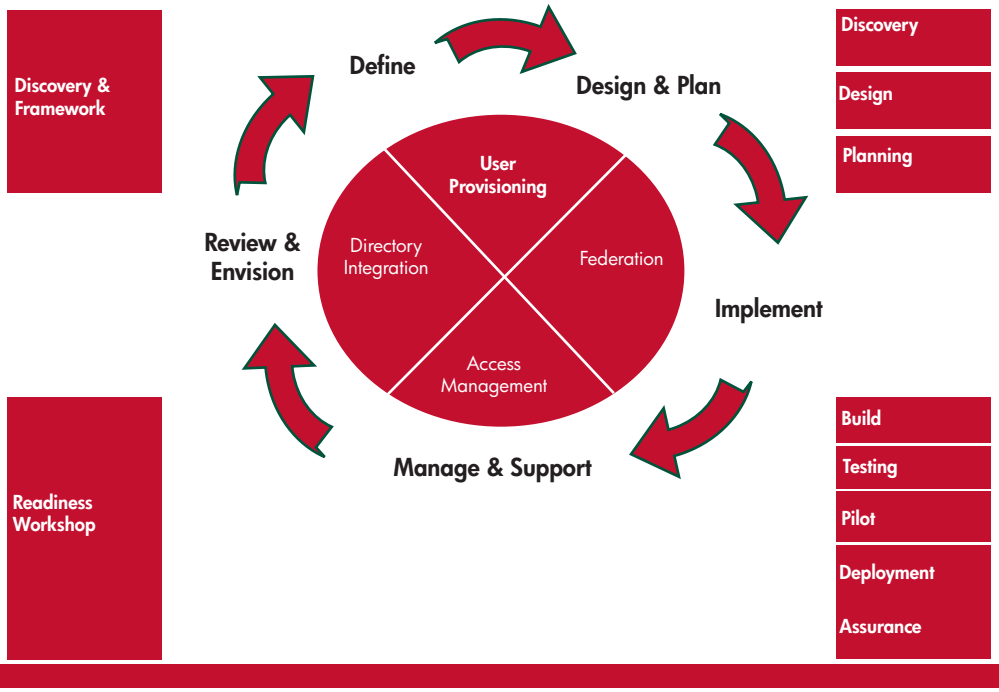

Design and Planning Services

Aligning with your architecture vision and roadmap, we develop detailed business, functional and technical requirements that will be transformed into the detailed design for the identity management solution, covering the functional solution areas of User Provisioning, Directory Integration, Auditing and Reporting and/or Access Management and Federation.

Identity Management Implementation Services

Using a proven HP methodology, our experienced consultants review, detail, and finalize your identity management implementation plan; conduct a production pilot; train and equip your deployment staff; fully execute the plan by deploying all solution components; conduct acceptance testing; and orient your operations staff for transition to production.

11. Identity Management Summary

Identity management is the ability to identify every user, application, or device across organizations and to provide flexible authentication, access control, and auditing while respecting privacy and regulatory controls. Delivered via a set of processes and tools for creating, maintaining, and terminating a digital identity, these tools allow administrators to manage large populations of users, applications, and systems quickly and easily. They allow selective assignment of roles and privileges, making it easier to comply with regulatory controls and contribute to privacy-sensitive access controls.

For HP, identity management is centered on the pervasiveness of identity management technologies and solutions:

- Identity management is about the management of user, application, and device identities.
- Identity management is about the management of identities in different contexts: enterprises, SMBs, consumers, and the public sector.
- Identity management deals with the management of the entire lifecycle of identities and their attributes.

HP considers privacy management, identity services, business-driven identity management, identity-capable platforms, and device-based identity management as important emerging identity management fields and has invested in specific research in these different areas that is driven from HP Labs. As an example of an end-to-end identity management system, the HP National Identity Solution provides governments with a high-performance, extremely secure, and extremely reliable credentialing solution. Similarly, HP can provide fully integrated end-to-end identity management solutions to meet any enterprise or public sector need.

Table 3-5
HP identity management solution offering summary

IDM Component	Solution	URL
Strong User and Device Authentication	Embedded Security for HP ProtectTools Smart Card Security for HP ProtectTools	www.hp.com/go/notebooks www.hp.com/go/desktops
SSO	Credential Manager for HP ProtectTools	www.hp.com/go/notebooks www.hp.com/go/desktops
AAA & Identity Repositories	HP-UX 11i Identity Management Solutions	www.hp.com/go/hpux11isecurity
Identity Management Services	IDM Consulting and Integration IDM Managed Services	www.hp.com/go/security Click HP security services link
Government	National Identity Solution (NIS)	http://government.hp.com

Chapter 4
Trusted Infrastructure

"Every change in the business triggers an IT event. If you get the infrastructure right, everything is possible."
-Bob Napier, late CIO, Hewlett-Packard

Running a business requires the availability and reliability of the IT infrastructure, which underlies most critical business processes. The reliability of the IT infrastructure is paramount. It implements the appropriate technologies to secure the end-to-end IT infrastructure, including data centers, networks, productivity tools, end-user desktops, and wireless devices. A trusted infrastructure and its network, host, storage, and print components form the basis of HP's security framework.

This chapter of the handbook discusses IT infrastructure security across networks, hosts, mass storage, and print infrastructure. It introduces the concepts related to trusted infrastructures, trusted computing, and directions in infrastructure technology. HP's security strategies for trusted infrastructure are also discussed, followed by detailed information about host, network, storage, and printing security.

1. Definition

Trusted infrastructures are composed of networks, hardware platforms, operating environments, and applications that behave in an expected and predictable way for their intended purpose. Trusted infrastructures support the IT applications underlying critical business processes. When IT infrastructure technologies fail to keep pace with emerging threats, we no longer trust them to sustain the applications we depend on in both business and society.

2. Purpose

The complexity of today's IT infrastructure exposes it to numerous threats. As shown in Figure 4-2, threats and challenges come from a wide variety of sources. These sources range from internal and external attacks to the risks introduced by common requirements for mobility, business partner connectivity, and outsourcing of IT services.

The need for a trusted IT infrastructure derives from an increasing reliance on IT systems to do everything from running businesses to running our society's utility infrastructures. Just as the dependence on IT permeates all aspects of society, security capabilities must permeate all aspects of the IT infrastructure. Security must be built-in, not bolted-on, at the platform level, at the network level, and in the very processes used for developing systems. A trusted infrastructure reliably manages and controls access to information assets while delivering the horsepower for critical business processes. It helps implement appropriate technologies to secure an organization's end-to-end IT infrastructure, worldwide.

Initially, security models in computing resembled a fortress with heavily guarded walls. As the power of computing, connectivity, and the Internet has become evident to businesses, this fortress model has manifested its limitations. The need for new IT security approaches has emerged as more companies harness the power of the network to do business online with customers and business partners.

Figure 4-1
Trusted Infrastructure

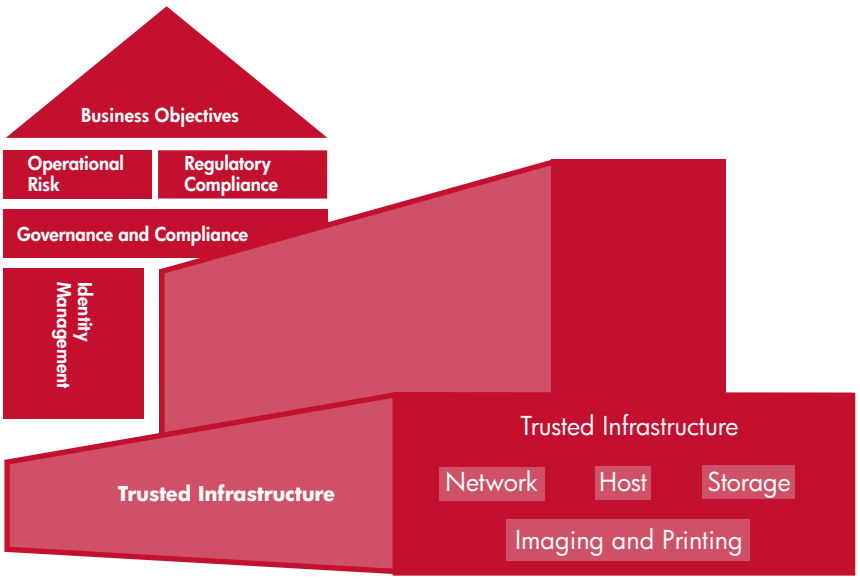

Figure 4-2
Challenges for IT infrastructures

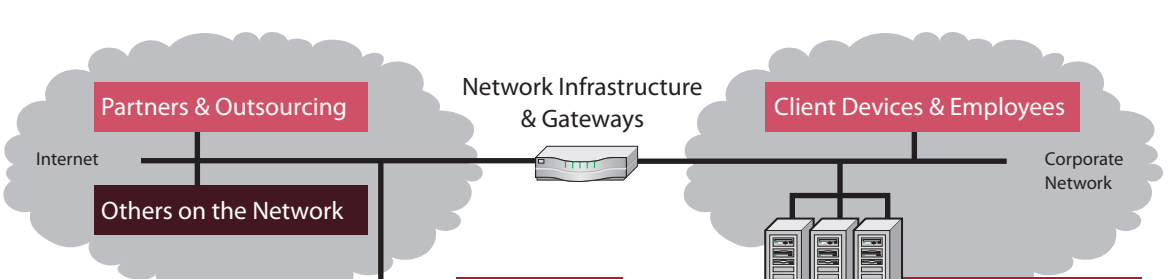

HP is working to meet trusted infrastructure needs across a unique breadth of products, from laptops and servers and storage and software to printers using HP network components. This is why HP created the Secure Advantage framework. Together with the HP ProtectTools client security portfolio and other HP security products and services explained in this handbook, HP Secure Advantage delivers a portfolio of products to meet customers' needs for secure data and infrastructure protection. Fortunately, HP has a long history in security and is leveraging this expertise to deliver the HP Secure Advantage portfolio. This is especially important today as customers adopt the 24-by-7 next-generation data center model that enables the shift of high-cost IT silos to low-cost, pooled IT assets in order to optimize infrastructures to reduce cost, increase agility, and improve quality of service. Security is a key enabler of HP's Adaptive Infrastructure (AI) offering that provides the platform for the next-generation data center and a linkage of security to other AI enablers such as IT systems and services, power and cooling, management, virtualization and automation.

Since security and compliance are an absolute necessity for businesses today; the HP Secure Advantage portfolio is designed to enable enterprises to fully automate, optimize and accelerate their IT infrastructures securely with proper validation, in order to achieve better business outcomes by mitigating risk.

2.1. Perimeter Security: Keep the Bad Guys Out

In the early days of computing, before its ubiquity in the commercial sector, dependency on IT infrastructures and the need for IT security were strongest for organizations in the government and military sectors. In these contexts, communication security and the need for access control to data were well understood and paramount. Conversely, the commercial sector perceived computing technology as a welcome performance and efficiency improvement - not a necessity. As a result, while computing technology became more widely available, technical developments motivated by the commercial sector focused on usability and performance.

For commercial computing applications, it was initially sufficient for organizational boundaries in the physical world to drive the requirements of mainstream IT security. The focus was on keeping the bad guys out using perimeter security. In the meantime, sensitive government and commercial organizations sought custom solutions to their IT security needs because they could not rely on off-the-shelf commercial technologies.

Many of the architectures that underlie major portions of today's infrastructures were designed to rely on perimeter security. However, as this perimeter security model has shown its limitations, so have the security models of the computers inside the perimeter.

2.2. Trusted Infrastructure: Let the Right People In and the Right Devices On

In today's world of remote workers, wireless users, trading partners, and connected customers, the expectations of perimeter defenses must be reexamined. Protecting the perimeter or point of contact with the Internet is still important, but it does not sufficiently provide end-to-end security. An effective security strategy must be far more flexible and sophisticated - simply posting a guard at the gate to the network is not enough. Infrastructure security requirements have evolved from keeping the bad guys out to letting the right people in. Legitimate users should have easy access to authorized resources, but they should be prevented from accessing unauthorized resources.

2.3. Ongoing Evolution

Organizations continue to use IT in new and changing ways. The evolution in computing and use models initiated by the Internet, connectivity, and mobility is still in its relative infancy. Modern businesses are interconnected with their customers and business partners, and they support an increasingly mobile workforce requiring seamless access to a company's IT networks from anywhere in the world. Similarly, as IT outsourcing offerings have become more comprehensive, more businesses are choosing a provider to host their IT systems. This recent and widespread evolution of how IT is used to run a business creates new challenges.

3. Infrastructure Technology Directions

An IT infrastructure is the collection of IT systems supporting a given set of IT applications. The core elements that comprise the fabric of a company's IT backbone are networking technologies and host technologies. Networking technologies incorporate hardware and infrastructure services and enable the secure and reliable transport of data. By contrast, host technologies incorporate hardware platforms and operating systems (OSs) and enable secure manipulation and storage of data.

Major developments in infrastructure technologies have occurred in the last several years. Mobility is a reality, networking is becoming pervasive, and IT infrastructures are becoming more adaptive and flexible at meeting business needs faster and at lower cost. Important areas of technological development in network security include network security architectures and network-enforced security compliance. Regarding host security, significant areas of development include OSs and hardware platforms.

3.1. Network Security Developments

3.1.1. From the Fortress Enterprise to the Adaptive Edge

As the enterprise IT infrastructure matures and adapts to new ways of doing businesses in an interconnected world, the network edge has moved outside the traditional physical enterprise. From the increasing pressure to provide mobility to a large part of the workforce, to the need for extending an IT infrastructure into a partner network to boost business efficiency, the edge of the enterprise IT infrastructure is less distinct. To adapt to a trend towards network commoditization, network architecture approaches need reconsideration. HP itself pioneered the "bubble architecture" in an attempt to apply good compartmentalization policy to network architecture and simplify security policy management.

HP promotes the compartmentalization philosophy with its Adaptive Network Architecture (ANA). HP uses ANA as a value-added differentiator for customers implementing network solutions such as IP communications, network consolidation and network security. ANA's goal is to logically compartmentalize an enterprise network based on the business needs of applications or hosted services and extend those compartments enterprise-wide, independent of geography, while enabling centralized policy management. The result is a simple network model that increases security and risk mitigation and reduces the cost of security policy management and controls while minimizing the time needed to implement change. Because ANA can establish network access and security rules for any given compartment upfront, adding systems to (or deleting them from) any compartment becomes easy as business needs change.

In this new environment, where the network edge reaches across public and untrusted networks into a remote host, client, or server, providing centralized administrative management of security policy all the way to the network edge is increasingly important. An organization cannot afford even one successful penetration of perimeter defenses; an attack jeopardizes the entire data network. To retain agility, businesses must manage the increased threat velocity and avoid ad hoc approaches. This creates a strong need for new approaches in data network architecture design that meet business agility needs while providing security with defense in depth. For example, the HP ProCurve Identity Driven Management solution addresses endpoint compliance requirements by allowing organizations to centrally control network access policy across wired and wireless access points (APs).

Self Test for a Trusted Infrastructure

1. Can I reliably distinguish a device that belongs to my organization's IT infrastructure from one that does not?

2. Can I tell that the firmware, software, and configuration of the devices inside my organization's IT infrastructure are in accordance with our IT security policies?

3. Can I trust the behavior of the platforms in my organization's IT infrastructure per our business objectives?

Conventional network perimeter defenses are challenged to simultaneously meet the needs of business agility and information asset protection. For example, firewalls are increasingly managed using exception lists, causing access holes within the firewalls and security and operational concerns. Another problematic trend is channeling a variety of application traffic over port 80 - the port commonly used for HTTP. Because port 80 is normally open to traffic even when the firewall is in its tightest state, this effectively circumvents firewall controls. In addition, some users provision direct external connections to support particular applications or projects, which can completely circumvent firewalls and other security controls.

3.1.2. Network-Enforced Security Compliance

Greater commonality in security functions across Local Area Networks (LANs), Wireless Local Area Networks (WLANs), and Wide Area Networks (WANs) is required. These three types of networks currently exhibit a large disparity in the level of security functions provided by their associated products.

To harmonize security enforcement is important to help maintain such security policies as access control across the network. It also facilitates central management of the entire infrastructure. Solutions exist today that help to implement such controls above a network infrastructure. However, additional efforts are necessary to provide holistic solutions that effectively deal with complex heterogeneous environments.

Pervasive and manageable security mechanisms are starting to be built into networks, with the help of standards such as IEEE 802.1x (for port-based access control). Additionally, infrastructure protocols such as 802.1x limit access to authenticated devices and users. When combined with a software solution for enforcement of end-point security compliance, these mechanisms help to support security policy decisions at the network edge. This permits such solutions as quarantining and remediation to take direct action on an authentication or compliance failure. Note that these approaches are not limited to the network edge - variations often can and should be used in the network core.

To better control connectivity to the infrastructure across a growing range of access technologies, new solutions are being developed in the field of Network Access Control (NAC). Initiatives such as Cisco's Network Admission Control (C-NAC) and Microsoft's Network Access Protection (NAP) are emerging as vendor specific solutions to this requirement.

While both initiatives provide a piece of the solution, organizations need an industry-standard solution to this industry-wide problem that manifests itself in the largest and most heterogeneous customer environments. To that effect, HP helped launch and promote the Trusted Computing Group (TCG) Trusted Network Connect (TNC) working group to produce specifications that support interoperability between individual vendor products across the network access ecosystem. The recent convergence of Microsoft NAP specifications with the TNC specifications in the TCG is an encouraging development in that direction.

Another important development in network security includes the emergence of behavioral approaches to mitigating threats. Building security features directly into the network creates proactive security management solutions. These types of solutions rely on cooperation from the components of the infrastructure for managing and mitigating fast-spreading threats. The Proactive Security Management chapter (Chapter 2) of this handbook provides details about these emerging solutions.

3.2. Host Security Developments

As discussed previously, most businesses now depend on a secure infrastructure. Yet they deployed platforms and OSs that were not necessarily designed with security requirements in mind; nor were they designed to work together well (if at all) in this regard. As a result, implementers of individual applications have been required to overcome these limitations and apply security protection themselves. A trusted infrastructure includes OSs and hardware platforms that offer reliability, manageability, and integration of security.

3.2.1. Operating Systems

When the necessary security mechanisms are built into the base of an OS, organizations can rely on standard enforcement mechanisms in the security architecture. Built-in mechanisms are harder to subvert. They also reduce dependence on correct implementation of the necessary security components in an application by (potentially) non-expert developers.

Security-relevant OS services include authentication, cryptographic libraries, intrusion detection, intrusion prevention, and compartmentalization technologies. When built into the core of the system, these security mechanisms are easier to control by policy, easier to manage across different OSs, and more reliably implemented by experts.

The release of the Microsoft Vista OS in 2006 demonstrates how OS security has become important in the mainstream and provides many security features built-in to this popular off the shelf client OS. Security features of Microsoft Vista will be discussed in more details in section 5.4.3.

3.2.2. Hardware Platforms

The utility computing platforms of the future provide virtualization of computing resources, such as CPUs, storage devices, and networks. These platforms, whether client or server systems, require integrated security mechanisms. For virtualization derived utility computing to succeed—from VMware to HP- UX, Microsoft Virtual Server, and Xen (an opensource virtual machine project)—businesses must be confident in the reliable separation and isolation of processes.

Modern platform and processor architectures, such as the IA-64 platform (Intel Itanium), are designed with security in mind. Other computing platforms in broad use today predate many of these security considerations. Most were initially designed with very different use models and functional requirements compared to today's expectations in typical IT deployments. For example, the original IBM PC, which is largely preserved in current PCs and mass-market servers, was not designed to meet the security requirements of present-day deployments. This is why HP has been leading industry efforts such as that of the Trusted Computing Group (TCG), to bring aspects of high-grade security technology to commercial IT systems at a low cost. Today TCG technology helps raises the bar for available off-the-shelf client and server technologies. More generally, HP aims to improve enterprise IT security by providing the foundation for enforceable security policies and strengthened identity mechanisms across a range of platforms.

Another important development is the emergence of virtualization technology in desktop PC and notebook devices, where it promises to offer better support for separation of duty between different processes running on the same machine. A recent example of commercial applications of virtualization technology on client devices is Intel's vPro. Moving forward, such technologies will help deliver enterprise manageability and security more reliably by isolating certain key functions from the operating system for remote management.

3.3. Encryption and Key Management Developments

The number of reported incidents of lost and stolen employee and customer data has risen exponentially over the past few years. There are many ways in which this data falls into the wrong hands, but lost and stolen tape drives and laptops containing critical personal data are some of the more common news stories available. This has lead to an increase in regulations such as the PCI ones (Payment Card Industry), to assure protection of consumer credit card data. The industry is focused on providing end-to-end solutions to implementing data protection, with a focus on enabling the use of encryption to protect data at rest across the enterprise infrastructure. This industry trend contributes to support HP's long-standing vision of the need for more trusted infrastructure where security is built in to the infrastructure components rather than bolted on as an afterthought. Protecting data at rest is one such capability that must become inherent to trusted infrastructure technologies.

Data encryption has been around for a long time, but was considered costly and hard to manage in the past. You can buy encryption engines to insert into the data path for encryption; however, these devices are specialized for specific use and each has their own key management capabilities. This increase in need for data protection has driven the encryption market to develop lower cost and more manageable encryption capability. For this reason, encryption is being built in to client, server, storage systems with eventual inclusion across the whole enterprise environment. The current trend is set on consolidating data-at-rest solutions to enable enterprise wide management of data encryption in heterogeneous infrastructures, but the need for end-to-end data protection solutions means that this trend is rapidly extending to broader data-in-motion and data-in-use models.

This trend has already started playing out with encryption solutions becoming increasingly mainstream across industry solutions, with data protection services built in to infrastructure platforms. In HP's portfolio this includes such solutions as Drive Encryption for HP ProtectTools, Microsoft Vista's BitLocker Drive Encryption, HP-UX's Encrypted Volume File System (EVFS), or the recently announced encryption industry standard for LTO4 tapes. And we expect to see more encryption capabilities being built into network switches and into individual drive mechanisms in the near future.

Coinciding with developments in increased encryption capabilities are improved key management solutions to assure ease of managing the encryption key lifecycle. Good key lifecycle management is imperative for encryption solutions; once something is encrypted it is not readable without the key. This also means that if the key is lost or corrupted, the data is non-retrievable.

As these technologies are built out, so will the standards definitions for encryption and key management. The myriad of various encryption and key management solutions leads to complexity and increased cost to manage these multiple environments. It is imperative that solutions be integrated around a common set of standards. This is also why various standards bodies such as OASIS, IEEE 1619.3, T10, TCG and others focus increasingly on encryption and key management.

4. HP's Strategic Focus

HP believes that security for trusted infrastructures must be built in and not bolted on as an afterthought. This belief requires a new level of maturity for IT security. Generally, IT solutions must provide improved mechanisms to underpin an organization's IT security policies, even in the face of developments such as utility computing, virtualization, and mobility.

4.1. Achieving Security through Open Standards

Creating trusted infrastructures using open and industry-standard technologies is central to meeting the real needs of IT managers. Open standards make it possible to provide security that is built in, manageable, and interoperable. The goal of enabling effective management of trusted infrastructures across large heterogeneous enterprises requires strong interoperability between vendor technologies, which demands collaboration and the development of (and adherence to) industry-wide specifications. For this reason, HP leads and participates in many standards bodies for infrastructure technologies. In fact, HP is an early founder and promoter of the Trusted Computing Group (TCG), created specifically to advance state-of-the-art technology in trusted infrastructures.

Interoperability is crucial for retaining business agility, particularly when businesses strive to achieve end-to-end security in a trusted infrastructure. HP's efforts within organizations such as the Internet Engineering Task Force (IETF) are aimed at creating the necessary interoperability interfaces.

Furthermore, efforts to advance the state of security mechanisms in the network must be combined with efforts to evolve device security. Trusted infrastructure solutions will rely on this. For example, approaches to network access control security will serve business needs only if they can be deployed in a truly adaptive and heterogeneous environment. They must interoperate smoothly and support the relevant industry standard(s).

4.2. Trusted Computing for Trusted Infrastructures

The TCG focuses on designing and standardizing security building blocks for the architecture of most types of computing platforms currently in use. This work supports the ability of those platforms to meet the growing need for more trusted infrastructure technologies. The Trusted Computing initiative collectively addresses new security requirements for computing platforms. At the same time, it preserves the openness and backward compatibility of platforms to remediate mainstream security holes and threats.

4.2.1. Trusted Computing Products

Trusted Computing security technology has become broadly available in business client PCs and notebooks from HP and other vendors, and in select server offerings.

HP has been a leader of the Trusted Computing initiative from the outset. HP's PC business and HP Labs teams include inventors and experts in Trusted Computing who spearheaded the Trusted Computing initiative many years ago. Trusted Computing is a great example of bringing "HP Invent" from HP Labs and forward- looking businesses into HP product lines and out to end users.

Trusted Computing delivers some significant benefits today, and some of these are manifested in HP products such as the HP ProtectTools Security Suite, or HP- UX Trusted Computing Services (TCS). While these products can be deployed today to better protect business applications and information assets, further advantages of Trusted Computing will increase when new hardware platform architectures are combined with redesigned OS software that can fully exploit the improved security attributes of the platform.

For more information about Trusted Computing technology, refer to the HP book “Trusted Computing: TCPA Technology in Context”, by Siani Pearson, et al., Prentice Hall PTR 2002, ISBN 0-13-009220-7.

TCPA and TCG History

TCPA stands for the Trusted Computing Platform Alliance. The TCPA was founded in 1999 by HP, Compaq, Intel Corporation, Microsoft Corporation, and IBM Corporation to address the issues of Trusted Computing Platforms. It is the predecessor organization of the Trusted Computing Group (TCG), which was established in 2003. The TCG was formed with a broader set of promoter members, a collection of typical industry consortium bylaws, and intellectual property agreements centered on reciprocal Reasonable and Non-Discriminatory (RAND) licensing. TCPA specifications in progress were brought into the TCG. The TCG extends Trusted Computing work to a broad set of platform technologies as well as infrastructure protocols and software stack interoperability.

Figure 4-3
Example of TPM internal architecture

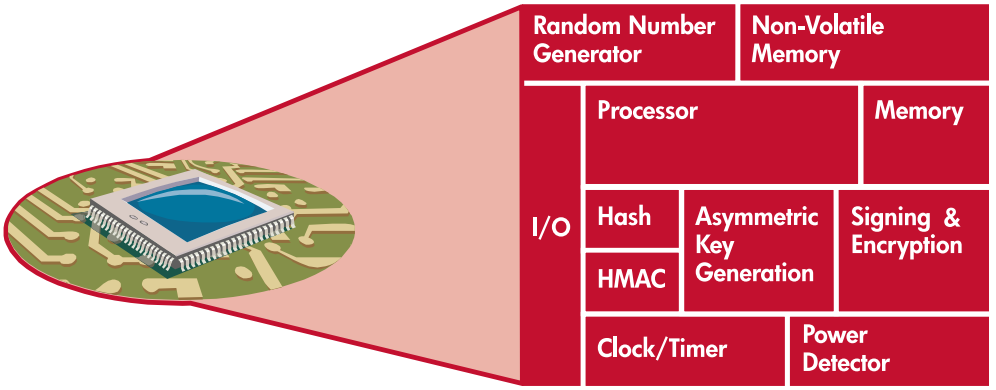

4.2.2. Trusted Computing Platform Functionality

The purpose of the TCG architecture is to prevent the subversion of key security features by software attacks on the platform. So that both local and remote users can trust reported information about the platform, it is necessary to protect the reporting mechanisms against software attack. The reason is obvious: the platform cannot reliably detect a software attack if its own software can be subverted. Protection against hardware attack is also necessary so that a remote user can trust reported information about a platform. This helps the remote user know, for example, that a local user has not physically tampered with the platform.

The Trusted Platform Module (TPM) is at the base of the trusted platform architecture. TPM is an enabling technology consisting of a dedicated security hardware device (with associated software) that meets TCG specifications. The TPM chip can be integrated with computer motherboards and many types of devices, including PDAs, notebooks, cellular phones, and servers. It provides multiple security functions including:

- Device authentication
- Attestation of software state on the platform
- Protection of secrets and stored data on the platform

The TPM includes all the functions that must be trusted in order for the TCG architecture to provide a set of features that cannot be subverted. Figure 4-3 illustrates the internal architecture of a TPM. From a business perspective, TPM-enabled devices create a way to manage business risk, manage assets, and protect critical infrastructures. For example, a TPM can support:

Data protection
- Stronger encryption
- Ease of use

Network protection
- Device authentication
- Protection of network credentials

Identity protection
- Strong, auditable, and attestable device identities rooted in hardware
- Built-in second-factor authentication methods for protecting a user identity

IT services and infrastructure for managing platforms
- OS-independent hardware-based policy enforcement on the use of and access to keys and data protected by the TPM
- Security policy compliance
- Software and hardware configuration management

4.2.2.1. Device Authentication

Trusted platforms provide mechanisms to help establish confidence in the behavior of a platform in an IT infrastructure. The basis for this confidence rests with the declarations from recognized and trusted third parties. These third parties can endorse a platform because they have assessed and measured the integrity of the platform. If the measurements meet a specific criteria, the third party states that the platform is trustworthy for certain purposes.

To associate trust with a specific platform, the trusted third party can certify a TPM cryptographic key. Two major classes of TPM cryptographic keys can reliably identify a trusted platform: non-migratable keys and migratable keys. Non-migratable keys are locked to the trusted platform from which they originate. By contrast, migratable keys can be moved from the originating trusted platform to another system, but only under the tight control of the owner of that system (the system user or possibly an IT administrator for that system).

Migratable and non-migratable keys exist so that a TPM can use them as cryptographic identifiers to prove that it deserves a third party's trust. For example, remote access software could use a TPM cryptographic key to uniquely identify the system in an IP security (IPsec) or an 802.1x authentication protocol to the back end of the IT infrastructure.

The concept of platform identity creates a reliable new security feature for the IT security tool kit: device authentication. Strong association of cryptographic credentials to a computing platform allows companies to personalize systems and issue credentials for recognition by the corporate network. Platform identity can also be used to configure the security credentials of a computing platform independent from OS security. This provides protection from mistakes or deliberate violation of certain security policies by OS administrators. For example, a security credential protected by TPM hardware can be controlled by specialized IT personnel (or a TPM administrator) to prevent copying or moving between machines by OS or domain administrators.

Features for device authentication are available today in a range of systems across the industry that have onboard TPMs. HP ProtectTools Embedded Security products provide features and mechanisms that can be used off the shelf for stronger hardware protection of user identity credentials, and they provide the building blocks to integrate more advanced physical device authentication to access control processes in IT infrastructure services.

4.2.2.2. Attestation of Software State

A computer platform has integrity if the OS and underlying firmware are tamper-free and applications running on the platform execute without interference. Existing security solutions assume the integrity of the platforms on which they operate. In particular, they assume that secrets can be safely stored and used on even the most open computing platforms, such as PCs.

Because platform owners are in control of their platforms' software environment and history (including interactions, physical modification, and software execution), owners may place trust in the integrity of their platforms. However, platforms are increasingly connected and exposed to threats from the Internet, which makes this confidence questionable. A third party is in an entirely different position than the owner, because the third party usually knows nothing about the environment and history of a remote platform. A third party, therefore, has no explicit confidence in the integrity of a remote platform.

For this purpose, the TCG defines an architecture that allows a computing platform to verifiably and reliably prove its integrity. This is achieved via a TPM-based mechanism that enables reporting of software and configuration measurements to a remote party. These integrity-reporting features are known as an attestation of the software state and configuration of a system. The features are not available today in mainstream platforms, because they are not integrated with mainstream OSs. The first commercial OS implementation that takes advantage of the TPM to verify software state is the Microsoft Windows Vista OS. The application that takes advantage of the TPM in this way is the BitLocker Drive Encryption feature, which protects data on the system. In the future, Linux systems, UNIX systems, and the next version of the Microsoft Windows OS are expected to take advantage of additional TPM mechanisms and support attestation features.

Attestation mechanisms provide the anchor for new architectures that will strongly rely on state information provided by remote systems. For example, a remote access solution could require systems that request network access to first attest that they have implemented the latest enterprise-approved security measures, such as anti-virus software and desktop firewall configuration on the client device. Another example is an information-flow security solution that controls access to and manipulation of enterprise data in an enforceable manner, based on security policy.

4.2.2.3. Protection of Secrets and Stored Data

On a trusted platform, a TPM provides logical and physical protection of secrets and logical protection of the data protected by those secrets. The TPM acts as a conventional cryptographic co-processor and its integrity-reporting mechanisms can prevent the release of secrets to inappropriate processing environments.

Specifically, a trusted platform provides hardware protection for keys and other secrets that typically encrypt files or authorize access to servers or other networks. The TPM can prevent the release of secrets conditioned upon presentation of a valid authorization value, the presence of a particular TPM, and/or the verification of a particular software state in the platform. This mechanism is known as the ability to seal storage to a given platform and/or a given software state on that platform. The TPM can therefore prevent inappropriate access to encrypted files and network resources that would otherwise be vulnerable to attacks, such as searching the contents of a hard disk, moving a hard disk to another platform, or loading software to snoop on other processes.

Because the TPM can enforce such policies, it is essentially a hardware-based policy enforcement mechanism for data decryption and cryptographic credentials.

Attestation of software state and sealed storage mechanisms will only be available to applications when OSs integrate the attestation features of the TCG architecture. As discussed earlier, this is beginning to occur for enterprise customers with the introduction of Microsoft Windows Vista and its BitLocker Drive Encryption feature. Today, HP ProtectTools Embedded Security products take advantage of the standard protected storage feature of a TPM to strengthen encryption solutions and provide a stronger tie between security credentials and a physical device.

4.2.3. Elements of a More Secure Platform

4.2.3.1. Embedded Security and TPM

In a PC, a TPM is attached to the low-pin-count (LPC) bus on the motherboard. A TPM provides mechanisms for root security functions and a hardware root of trust in support of OS security. Beyond providing well-understood cryptographic functions, TPM features support the design of new OS architectures that create a chain of trust, which is built from the TPM hardware root of trust and extends to software on the platform. With a TPM, a typical chain of trust can provide strong cryptographic attestation (across a network) of the state of a local platform's firmware, hardware configuration, OS, and software configuration. Combining a TPM with higher-level software creates the basis for strong, hardware-based policy enforcement for the first time in mass-market systems. HP workstations, desktop PCs, and notebooks are available with a TPM.

4.2.3.2. Operating Systems

OS support is expected to gate the most widespread commercial availability of the Trusted Computing platforms. Those platforms will integrate TPM features and combine other components - such as new CPU and chipset security architectures from AMD or Intel (for example, Intel's TXT Technology) to provide security mechanisms that directly benefit higher-level applications.

As noted previously, Microsoft is expected to build on these technology components in future versions of the Microsoft Windows OS and provide remote software state attestation features enabled by the Trusted Computing architecture. Linux and UNIX vendors are expected to make use of these technologies in the same timeframe to create similar capabilities for these platforms.

4.2.3.3. Applications

HP's ProtectTools products enable legacy applications to take advantage of the TPM transparently through standard interfaces such as the Microsoft Cryptographic API (MSCAPI) and the Public-Key Cryptography Standards (PKCS) #11 interface. They also provide applications designed to use a TPM to enhance data security. Newly developed applications will use TPMs on computing platforms. OSs that build a chain of trust from the TPM will also provide benefits for the management of trusted infrastructures, independent of individual applications.

4.2.4. Trusted Computing Across the Infrastructure

The benefits of Trusted Computing are available to virtually any device that contains a processor and an OS/environment, runs applications, and communicates with other devices via networks. The value of this emerging technology becomes greater in more open platforms; it helps attest to appropriate state and configuration without restricting and locking the platform completely. The expectation is that Trusted Computing will appear in all relevant form factors over time, including PDAs, servers, and mobile phones.

4.2.5. Security and Privacy Issues

Not surprisingly, many security-enhancing technologies have privacy implications. Privacy requirements are dependent upon the context in which they are viewed. In some transaction usage models, increasing the security of data requires identification from an actor (user or system) that wants to access the data. In other models, anonymity helps ensure the security of the actor's identity or Personally Identifiable Information (PII).

In the past, Trusted Computing was mischaracterized as a privacy threat. In fact, Trusted Computing specifications have been developed with specific, privacy-sensitive principles to allow for secure IT solutions that respect privacy. Trusted Computing contains building blocks that, used correctly, can protect the privacy of data or the actors wanting data access. Notably, the Trusted Computing specifications have consistently built privacy considerations into the design of the technical architecture. Various mechanisms support the protection of private data and avoid approaches that create privacy concerns (such as a visible, single, and unique identifier for a platform). From mechanisms that support the creation of pseudonymous identifiers to designs that let platform owners opt - in to use the technology, the technical specifications carefully consider the protection of PII.

In the past, Trusted Computing was mischaracterized as a privacy threat. In fact, Trusted Computing specifications have been developed with specific, privacy-sensitive principles to allow for secure IT solutions that respect privacy. Trusted Computing contains building blocks that, used correctly, can protect the privacy of data or the actors wanting data access.

Trusted Computing can be effectively deployed across a variety of use models with differing privacy attributes. This includes meeting the strictest privacy legislation and providing the basis of privacy-enhancing technologies for future IT solutions and platforms. HP's ProtectTools Embedded Security products comply with the privacy-sensitive and user-control spirit of the TCG specifications. In addition, HP Labs is actively pursuing the design of new privacy-enhancing applications of Trusted Computing with the broader research community.

The TCG has active working groups to address the different types of platforms. The working groups are designing specifications for the Trusted Computing architectures of the various device categories and their use models. All of the TCG's work focuses on manufacturer - and vendor-independent specifications to enable interoperability of implementations.

In addition, the TCG is focusing on infrastructure protocols and mechanisms to design interoperability and support for new trusted computing features. HP leads and participates in these efforts, including the TNC working group. This work will allow the next generation of infrastructure services to seamlessly use Trusted Computing technology across multi-vendor platforms, OSs, and applications, supporting the design and deployment of truly heterogeneous trusted infrastructures.

4.3. Network Access Control (NAC)

NAC has evolved over a number of years to become viewed as a critical part of layered defense or defense in depth network security. However there is still some lack of consensus around what NAC really means. This is due to a combination of early proprietary approaches, Cisco's Network Access Control scheme (C-NAC) Microsoft's approach with Network Access Protection (NAP), and an early lack of standards. HP has worked to address the lack of standards through the TCG with TNC (Trusted Network Connect). HP sees NAC as a standard part of network security offerings over the next few years, and it is critical that a common understanding of core NAC capabilities, essential services and common standards is established.

4.3.1.What is Network Access Control (NAC)?

HP defines NAC as the combination of software, hardware, services and processes designed to protect a network from untrusted or unsecured endpoints. NAC is primarily a network security element, intended to protect the network and its resources from harmful users and systems/devices. NAC controls and restricts access to network resources based on certain criteria and business policies. In its most basic form, NAC allows a network administrator to restrict network access to authorized users and/or devices. However, many organizations have the need to provide, or can benefit from providing, different levels of access depending on the role of the user. For example, employees have access to internal network resources and the Internet while guest users are only provided access to the external Internet.

The need for protection from malicious software is accomplished by evaluating the "health" or "security" posture of devices connecting to the network. The required posture is defined by organizational policies and is based on checking for things such as operating system version, patches, security software (anti-virus, anti-spam, firewalls, etc.), security settings on common software installations, or other required or prohibited software. NAC goals can be further complicated by the fact that today's network is often comprised of network access requests from devices that are not under direct organizational control, such as contractor and guest laptops. Furthermore, the need to understand and comply with regulatory agencies and company policies alike drives a need for the organization to seek solutions that meet this goal, often with fewer resources than ever before.

NAC must deliver on securing the organization based on the organizations governance model and business objectives. This means that NAC should not become a barrier to business and must be able to handle exceptions or emergencies as well.

4.3.2. NAC Benefits

The business benefits of proper NAC solutions are significant, and include:

• **Improved governance and compliance**: When dealing with regulatory or corporate compliance requirements, NAC allows an organization to significantly improve their ability to ensure that access to specific systems and data is only available to specific authorized devices and users that comply with policy. Additionally, with the right implementation, the ability to audit and report on the environment is increased. NAC implementations then allow for the high level governance capabilities to be aligned with common network security due diligence used in many different governance frameworks.

• **Improved security posture**: NAC provides an additional protection layer for an organization's Defense in Depth or Layered Security requirements. While it requires analysis specific to an organization, the goal is to minimize risk to the network business resources from unauthorized, unhealthy and out-of-compliance devices and endpoints. By doing this, NAC can reduce unnecessary exposure of corporate assets; for example, if a PC is running peer-to-peer software, then there is a risk that confidential documents could be inadvertently shared. The presence of such software could be detected, audited and acted upon to ensure that the PC does not get onto the secured network with such software in operation.

• **Improved operational cost management:** Organizations face tremendous pressures to prevent breaches, ranging from virus infection through data loss, while at the same time maintain or decrease cost structures. Investing in NAC capabilities allows an organization to increase the security posture while ensuring that fewer issues need to be resolved post-breach. Costs associated with resolving security breaches after the fact are often hard to quantify, but a lot of data is available to evaluate the likelihood of such events.

4.3.3. NAC Challenges

NAC solutions today are not overtly complex in their goals or implementations, but can also be considered relatively simplistic in their enforcement capabilities. The largest challenges facing NAC today are:

• **Politics**: Like many technologies, NAC has the potential to significantly change the way in which people will need to work when using networked resources. Initial implementations can fail if they create too complex remediation processes, or worse, force a user into a dead end where they are unable to work at all. The commonplace example is a critical deal being lost because some individual could not get on the network to obtain or submit critical time-sensitive information. Make that person an executive and the example can often become more serious.

• **Complex integrations**: In order to successfully deliver NAC, it is required that all parties work well together. Many vendors provide their own partner integration programs.

• **Legacy or limited endpoint capabilities**: While endpoints such as PCs and servers can run agents or respond to remote queries to determine their software posture, devices such as networked printers, phones, PDAs and so forth usually do not yet have the standard capability to respond to standard or even alternate NAC challenges, such as web access redirection or 802.1x-based authentication. Therefore organizations implementing NAC usually end up using exceptions such as MAC or IP address authentication, or implementing guest VLANs. Since MAC or IP address authentication can often be spoofed, it is important to consider carefully the security implication on a NAC deployment, and implement a separate guest VLAN when possible. tools.

• **Proprietary solutions**: Most vendors have their own agent technology. Firstly the initial lack of common baseline functionality and standards has forced vendors to implement or OEM client agents that cannot work with other solutions. An ongoing disconnect between standards and proprietary solutions remains at the network level, which limits comprehensive innovation across the NAC management space, in terms of standard integrations with tools such as SIM/SEM, change management, network management, and similar tools.

HP is working on all these areas through a combination of standards activities, partner integrations and advanced service delivery capabilities. In addition, HP ProCurve's unique identity and immunity solutions already provide advanced NAC capabilities across the network, to the network ports and endpoints that are part of the evolving NAC environment.

4.3.4. NAC Futures

To address many of the challenges associated with NAC, HP sees the NAC market evolving to deliver the following:

• **Standardized NAC infrastructure**: With the work of the TCG's TNC working group, HP believes that standards for NAC infrastructure will help meet customer needs for interoperability between NAC level products. For example, Microsoft's recent NAP alignment with TNC will have a significant impact on creating a common NAC framework. HP will continue to work with vendors and standards bodies to deliver a standardized NAC infrastructure. Further, the increased use of interoperability testing will ensure that the infrastructures will provide for easier deployments.

• **Device Identities**: HP sees the need for secure device identities to be implemented to support NAC security architectures. Using existing standards such as TCG's TPM specifications, the TCG's TNC working group and the IEEE 802.1AR work will better address network infrastructure security needs: endpoints will be able to provide stronger security assurances with hardware protected device identity credentials, and signed health statements to a NAC ecosystem.

• **Standardized NAC integrations**: To minimize friction between governance models and network security initiatives, it is critical that NAC be able to support and respond to an organization's supporting Security Information/Event Management (SIM/SEM), change management, network management, and similar tools.

• **Behavior-based NAC**: Linking NAC implementations with network monitoring capabilities allows for legacy devices to participate more fully in a complete NAC environment, while appropriately mitigating the risks associated with their lack of NAC device client capabilities. This will evolve into a cyclical relationship between these solution areas delivered by standardized NAC integrations.

• **Virtualization and hypervisor evolutions**: With the emergence of virtualization technology on endpoints, we expect to see the development of hypervisor-level NAC solutions for endpoint compliance enforcement. Proprietary technologies such as Intel vPro are beginning to take advantage of hypervisor technology to isolate and secure network security policy enforcement on individual endpoints, and we expect such implementations to integrate with NAC architectures moving forward.

4.4. Secure Development

The root cause of most security incidents (beyond the perpetrator of an attack) is typically the exploitation of a vulnerability that allows the unintended outcome. Of course, people and processes can create significant vulnerabilities, and there are many ways to track known vulnerabilities, patch them, and block them. However, this reactive approach is not sufficient by itself. To be truly proactive about dealing with security-related vulnerabilities, the responsibility shifts upstream in the development cycles to the development teams who create the software or firmware in the first place. That's where the vulnerabilities are created unknowingly. It is clear that developers must be more aware of best practices and bad practices to create less vulnerable products and solutions. This is the motivation behind HP's secure development initiatives. For HP, secure development is an ongoing process that begins with awareness and education and continues all the way through the product lifecycle. This is how HP works to produce secure and trusted products and solutions minimizing bugs and flaws that have security implications as well as building in security right from the start.

4.4.1. Minimizing Flaws

As evidence of the importance of secure development practices, various worms such as Code Red, Nimda, Blaster, Slammer, and Sasser have caused havoc on public networks, private networks, and home systems. The root causes of the vulnerabilities - which were exploited by these attacks-comes down to a single untrusted library call, a failure to prevent a memory structure from overflowing, or some other insecure software development practice. Using such attack analysis information, HP's secure development initiatives are aimed at minimizing known bugs and flaws that have security implications through education for the developers and internal tools/methodologies. In addition, developers can add security technology to design and solution architectures. The combination of secure development methods, internal tools and the inclusion of security technologies early on or upstream in our product lifecycles means that HP can increase solution quality and trustworthiness without significantly impacting the product's time to market or costs. For our customers, this means better overall quality, value and cost savings by avoiding security issues in deployment.

4.4.2. Developer Education

HP has a worldwide security education program targeted at all internal developers. Both general developers and security-focused developers need to learn how to make less vulnerable products by minimizing bugs and flaws with security implications. The curriculum also includes courses for security-focused developers who also need to learn how to incorporate specific built-in security technologies. The program includes best-practice white papers, on-demand seminars, computer-based training modules, and instructor-led courses.

4.4.3. Product Development Lifecycle

Other parts of HP's secure development initiatives focus on constantly improving product development lifecycles. HP has added security-focused steps to each stage of product lifecycles. For example, risk assessment and vulnerability assessment techniques are used during the design phases, and the testing phases present the opportunity to perform both component- and system-level security testing. Processes and methodologies are brought into HP development lifecycles, along with source code, application- and system-level vulnerability scanners, and threat assessment tools. HP uses both its own HP-invented tools and best-in-class tools from third-party vendors.

4.4.4. Vulnerability Analysis

As mentioned above, vulnerability assessment tools or scanners play an important role in the product lifecycle and can serve both proactive and reactive functions in the development and maintenance of secure programs. Proactive vulnerability analysis refers to employing risk assessment and vulnerability tools early on in the development process. During architecture definition, risk assessment can highlight the areas of highest risk, assert security best practices to secure those high risk areas and even guide the application of vulnerability assessment tools based on business requirements and risk prioritization. Reactively, in later stages of product development and testing, assessment tools can be run on the running programs, applications and the source code itself.

This provides a quality checkpoint or testing phase that can catch security bugs or flaws that have been created. Further, most assessment tools provide corrective information about how to resolve the bug or flaw. Vulnerability assessment tools look for known patterns of bugs and flaws, and their databases of these now exceed 4,000 known problems today, and grow frequently! But even more advances have given us the ability to look for unknown flaws or bugs that lead to security problems - these include techniques such as dynamic attack methodologies and/or fuzzing techniques to reveal previously undetected flaws in an application. These are flaw-finding techniques where a testing tool (fuzzer or dynamic attack generator) sends random input to the program being tested, looking for input that can lead to an exception, crash or server error in the case of web apps. The testing for unknown and known flaws occurs simultaneously or as part of the same process for most modern scanners, such as HP WebInspect.

The proactive approach for applications is to perform such vulnerability analysis of applications while they are under development and to perform application security assessments against deployed production applications on a regular basis to ensure the security state of the application remains known and good in the face of constantly evolving threats. Interestingly, compliance issues are driving more vulnerability assessment directly. For example, the Payment Card Industry Data Security Stander (PCI DSS) requires application assessments of applications in development and on a regular basis after deployment. If you process credit card transactions online, you are required to comply with PCI DSS and therefore you are required to use vulnerability assessment techniques.

All the above effort results in products, solutions, and services that are built with fewer bugs and flaws (with security implications) and designed with security in mind. In addition, HP sells the assessment tools to build your own secure development processes, and HP Services is making secure development practices and expertise available directly to customers. Secure development services from HP include education and training as well as threat and vulnerability assessment.

5. Host Security

In the book *The Mezonic Agenda*, a respected security expert, Chad Davis, chases an international conspiracy to sway the U.S. presidential election. Early in the story, Chad avenges an embarrassing situation by quickly updating a very visible web page - in spite of excellent client-side security. Specifically, he discovers that JavaScript validation routines in the client HTML are the sole bastion against an SQL-injection attack. He ultimately pegs down his arrogant colleague by entering:

```
"; EXEC master..xp_cmdshell 'echo I am
insecure! >?c:\inetpub\wwwroot\home.html"
```

There are two very important - and relevant - lessons to learn from this example:

• Security is no longer about perimeter defenses. In the early days of mainframes, there were single points of access to computing resources and data that were easy to secure and manage. Now, however, data and processing power is distributed throughout the organization, in hundreds of different servers covering thousands or even millions of clients. There is no good place to draw the perimeter, because the network topology is so dynamic that it is generally impossible to enumerate or calculate.

• Everybody forgets the server. Even well-trained and experienced security architects adopt the perimeter model and neglect to recognize the inherent vulnerabilities in application and OS code. They may also believe that platform security is not cost-effective because, unlike most other countermeasures, it requires frequent administration and introduces significant complexity into both security policy planning and security administration.

What is the point of developing a secure network perimeter with layered levels of firewalls, strong passwords, and intrusion prevention and detection systems, if a simple buffer overflow exploit opens the web server to unfettered access by unauthorized users?

Hardened OSs have historically been hard to use, hard to integrate with the environment, and difficult to verify as secure. The most important question is whether platform security yields the expected returns. Is the effort required and Total Cost of Ownership (TCO) too high relative to the estimated threats and the value of the assets being protected? Although platform security can be very effective, it may not always be worth the cost.

HP looked intently at this issue and enhanced the delivery of platform security through the operating environment. The result is new tools and techniques that reduce risk to the enterprise without ballooning TCO or creating an unacceptable customer/user experience. This section shares our view of platform security and its value to the overall infrastructure.

5.1. Environment

Organizations in specific sectors and industries - such as financial services, the military, and the intelligence community - have used strong platform security for decades. In the majority of other sectors or industries, the reaction to hardened OS products was a question of why an organization would want to implement them given the additional effort and inconvenience required to achieve the high level of security. This question illustrates the fact that security and platform security were not high enough priorities for IT administrators to warrant the cost and additional effort. In addition, as recently as ten years ago the platform security community did not differentiate between life-and-death and profit-and-loss market segments. Instead, it offered a single solution to satisfy all higher security requirements.

5.1.1. Battlefield Protection, Enterprise Overkill

Legacy host security grew out of Cold War technology and thinking. Early offerings were based on the Compartmented Mode Workstation, an intelligence desktop that allowed secure assembly of field data gathered from spies all over the world. Correctly implemented, even the spies did not know what the data meant, because they never saw it all pieced together - hence the compartmented approach.
In the context of international espionage, host security was defined as:

• Separation or compartmentalization of different kinds of information
• Separation of powers or authorizations, so that nobody had all the keys
• Separation of various activities into individual tasks, each with its own associated privilege
• Highly granular accounting systems or auditing that tracked each user and system event

The main idea was compartmentalization or layering, much like the watertight compartments in a submarine: a failure in one section did not flood the entire submarine. However, transferring this security model to the connected enterprise raises some issues.

5.1.1.1. Drawbacks of the Legacy Approach to Host Security

This approach to host security usually has a high TCO for several reasons:

• Systems are designed for split administration (prohibiting one person from managing the whole system), which means higher personnel costs.
• The level of security usually adopts a firewall mentality: what is not expressly permitted is prohibited. That model works well for routers sorting subnets, but it breaks quickly for complex applications trying to communicate with scores of OS services over dozens of interprocess communication (IPC) connections.
• Secure platforms are often simple, functioning without today's modern operating environment management systems. As a result, implementing routine functions requires extra effort.
• Secure platforms require custom, security-aware applications that are specifically written to behave in a way acceptable to a completely hardened OS.

5.1.1.2. Benefits

The infrastructure environment is not perfect. Ever since the early days of computing, programs have had bugs. Every piece of software contains some level of program faults, design mistakes, partially implemented features, and possible holes that developers may have been aware of but considered safe from exploit.

Layered security acknowledges upfront that systems and software always contain defects or bugs of some kind. Under the right set of circumstances they will eventually break or be compromised. Like the proverbial submarine, military security is equipped with watertight doors, with the full expectation that one or more layers will not withstand every attack. A hardened operating environment is one of the most effective ways to prevent broken and rogue applications from violating OS security policies.

Figure 4-4 illustrates how attacks can penetrate layers of defense. If A1, A2, and A3 are attacks using different exploits, each attack is stopped at different layers in the diagram. The layered host security model assumes that bulletproof protection does not exist, and that some number of attacks will be successful. The goals are to set up several layers to prevent as many attacks as possible, survive the attacks that do occur, quickly regroup, assess and repair the damage (as possible), and continue operating as planned. Application software is often developed with the assumption that protection comes from somewhere else. It relies on the operating environment for protection from other users/processes on the system as well as the external, distributed, and networked environment.

Figure 4-4
Host security layers

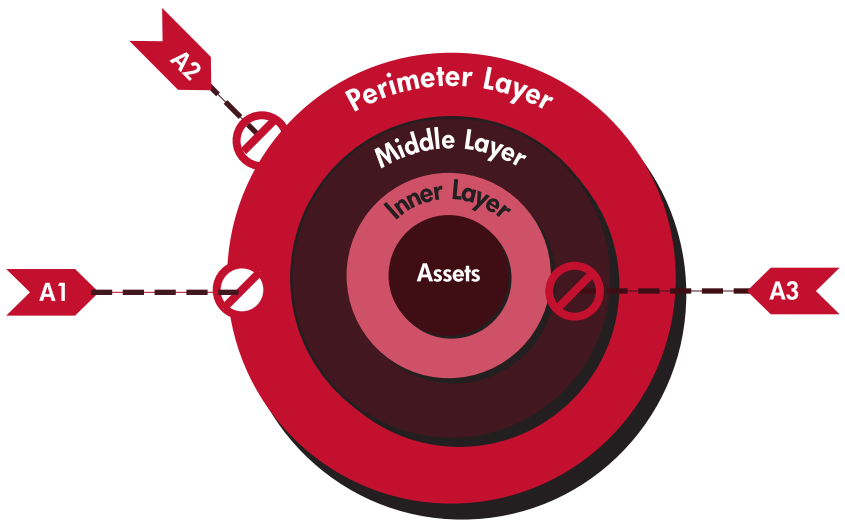

5.1.2. Commercialized Forms of Military Systems

Several attempts have been made to transition the military approach to platform security into the enterprise space. Most of these either simplify the compartment layout (preconfigured systems) or accept the requirement for ongoing, highly specialized consulting. Whether using preconfigured or customized systems, the administration costs are relatively high compared to commercial systems, meaning the TCO comparison has not traditionally been favorable. In addition, security is sometimes so tight that some applications or services simply cannot run on the transitioned systems, regardless of expenditure.

The emerging requirement expressed by HP customers is for strong host security with lower TCO, coupled with the flexibility to accommodate a broad range of applications, platforms, management tools, and markets. This led HP to analyze the host security market, along with some of our other markets, to find out whether high security can be simple, available, and cost-effective. The catch for the high-security operating environment is usually the TCO, which is generally very high when the asset value is high. However, this can be overcome if the host is chosen well, matched to the enterprise infrastructure, and surrounded by properly crafted, implemented, and enforced security policy.

5.2. Principles of Design for the Enterprise

As HP has considered how to make high security produce a higher return over a broader target market, we have made several discoveries. First, life and death situations merit a cost/benefit analysis that is different from normal business environments. The military model addresses non-financial losses, such as the loss of human life and the collapse of governments - events much more catastrophic and untenable than a simple reduction in profits. When military models are considered for business purposes, they are usually out of balance on the expense side. In other words, military security, quite appropriately, is not intended to produce a financial return on investment (ROI).

Second, these layered approaches generally assume that any manageable level of complexity is acceptable, even if the administration effort is high. This may be appropriate in a life-or-death environment. Without the military/political threat, however, the ROI is not justified, except for a few highly sensitive niche markets such as high-end financial services. Armed with these observations, HP set out to determine whether the high-security approach can be retooled to accommodate an enterprise environment. We began to ask detailed questions about high security, exploring different aspects of military and layered security models and ways to implement these features without placing the TCO equation out of balance.

5.2.1. Easily Administered Layers

Regardless of how high security is implemented, layering is still needed because it is fundamental to containment and risk mitigation. It must be designed into the core operating environment to minimize tampering. Luckily, the internal mechanics are not the problem. The real issues with the military model are configuration and administration. In a typical model, compartments are defined by many layers of indirection, which leads to complexity and lack of flexibility.

In an enterprise environment, the average user depends on a specialist to define layers in most high-security platform systems. Platform layers must be simple to configure and maintain. In fact, a layered model that is both role-based and rule-based eliminates most of these administration issues, allowing the platform administrator to easily create and change configurations. If a system sets up the layering correctly and the administrator describes the layers in a straightforward way, administrative costs drop dramatically - contributing to a significant increase in ROI.

5.2.2. Flexible Role-based Access Controls

It is clear that several other military security features have value in the commercial space. For example, the root account on most systems is all-powerful and able to execute any system command at any time. Even experienced administrators use root only when they require elevated privilege. This observation drives the concept of designing tools that adopt a valid role for specific commands, only for the time required to perform a specific task or function.

Administrators gain privilege for each command that needs it, for the time required, and in a specific area of responsibility. These Role Based Access Controls (RBACs) can be managed on a user-by-user basis. This permits specific users to assume more powerful roles or privileges, depending on their job requirements.

5.2.3. Realistic Privilege Allocation and Management

Another aspect of platform security is the level of privilege assigned to system capabilities. In a true military environment, every system function has its own privilege level - much like putting lockboxes inside a locked desk drawer, inside a locked office, and so on. In the enterprise environment, there are two problems with numerous locks. First, far too many checkpoints are required for a relatively simple system operation, for example, printing a document. The extra checkpoints cause normal system operations to be slower and make applications more complex. Second, the privileges often overlap or create unnecessary redundancy.

Clearly, privileges have value in a secure environment. For example, it is valuable to control the ability to erase a disk drive or transmit files over the Internet. But the privileges must be assigned to enterprise-level activities, such as erasing a disk drive or sending files, rather than to the minute collection of system operations that make up these activities. If platform privileges are reallocated to a higher level of abstraction, they provide useful protection without incurring unnecessary costs - and thus lower the TCO.

5.2.4. Balanced Security and Performance

HP realized that the military model focused very little, if any, on performance. If the system did not run fast enough, more powerful hardware could be obtained. This is a critical difference between a life-or-death decision and an enterprise's profit-and-loss decision. Espionage and battlefield situations usually involve escalation of force, little or no consideration of cost, and more and bigger hardware.

In an enterprise environment, bigger is not necessarily better or more cost-effective. In fact, enterprises tend to accept a slightly higher level of risk in order to reduce costs or raise ROI. And if the security is particularly demanding of resources or effort, an enterprise might disable security features, which may be an appropriate choice for the environment. Hence, there is a need to link platform security with system performance management tools. For example, the HP-UX 11i v2 security containment and processor partitioning solutions known as vPARs (virtual partitions) and nPARs (node partitions) tie into workload management, process resource management, and the HP ServiceGuard product. (For more information, see www.hp.com/go/unix).

In order to increase performance while maintaining security, other design goals emerge such as keeping applications in their designated compartment and preventing them from using more resources than appropriate. Another example comes from the virtual partitioning architectures: when an application needs more computing resources, it must be able to automatically add resources without compromising security. These examples of combining performance with security goals illustrate how the role of security is to support new models of operation, not to administer security for security's sake.

5.3. Implementing Secure Platforms

Secure platforms are, to a great extent, constructed during implementation and integration. They are building blocks or foundational elements. Although not all secure platforms involve changes to the OS, most of them are so tightly integrated with the OS kernel and other core operating environment functions that it is unreasonable to design a platform security system in the field. Based on that understanding, two conclusions can be drawn:

- The selection, configuration, and implementation of a solution is more important than the availability of specific security features. In other words, the security of the platform depends greatly on how it is configured and implemented.
- Because platform security architecture is largely predesigned and made configurable, a good platform security implementation should place increased emphasis on security management. A secure platform is not effective unless it is accompanied by solid security policy that supports and surrounds the platform.

This section briefly reexamines some of the fundamentals of security architecture, focusing on how they relate to implementing and ensuring a secure platform.

5.3.1. Security Architecture Models

Unfortunately, the first step in constructing a security architecture model requires abandoning the secure perimeter model. The distributed nature of computing systems makes the perimeter difficult to locate and secure. Examining the basics of security modeling helps to understand why perimeter-based thinking is flawed. Security modeling is the fundamental baseline for security assurance, that is, for assessing and verifying the security of a given implementation. There are probably as many security models as there are ranking experts. Examining a few of the most common models illustrates what they have in common and why the perimeter model breaks down in enterprise computing.

5.3.1.1. State Machine

State machine is the core of most security modeling and verification systems. In a state machine model, the world is divided into subjects and objects. Subjects do the acting and objects are acted upon. Each subject (program, process) and object (file, memory range) is assumed to have states, which change over time (state transitions).

A simple example of subjects and objects might be {man, boy, bat, ball}. Acceptable states might be {accelerating, decelerating, stationary}. Most of the acceptable state transitions would involve the boy accelerating the ball with the bat in such a way that the ball does not use the man as the unexpected subject of a deceleration state change. In a secure system, the goal is to ensure that every possible state change or state transition is considered to answer the question: if the system starts in a secure state, are there any actions of subjects on objects (state transitions) that can cause the system to become insecure?

5.3.1.2. Bell-LaPadula

The Bell-LaPadula model, dating from the 1970's, mirrors the classification system used by most governments to label sensitive documents. The fundamental principle of Bell-LaPadula is the way it imposes a lattice or hierarchy of subjects and objects. It facilitates a quick comparison to decide whether a given subject is allowed to perform a certain action on a given object. It hinges on proper labeling of subjects and objects, and the discussion of levels, labels, domains, and dominance can be very complex. The Bell-LaPadula model is concerned with the confidentiality of data.

5.3.1.3. Biba

Essentially, the Biba model is identical to the Bell-LaPadula model, except that it deals with data integrity. A user may be authorized to access certain data, but how does the user know that it is the right data and that it has not been corrupted? This model also makes use of subjects and objects.

5.3.1.4. Clark-Wilson

Clark-Wilson is a proxy-based integrity model, stilted toward the commercial environment and focused on separation of powers or authorizations. The goal is to prevent authorized users from making unauthorized changes to information.

5.3.2. The Trusted Computing Base and Dynamic Proliferation Model

Loosely described, the security perimeter is equivalent to the trusted computing base (TCB). The TCB is roughly defined as the set of subjects and objects over which the security administrator can have reasonable control and assurance. The components that can be cleanly identified, mapped, analyzed, and controlled by the security administrator fall within the TCB and earn certain levels of trust. Things that are not as neatly managed fall outside the TCB and cannot be trusted (within the confines of this model). In a perfect world, the security perimeter includes all enterprise data, users, and resources and an appropriate (reasonable or cost-effective) level of trust through various security policies and controls. Even things coming in from outside the TCB, such as network connections or anonymous customers, can be identified in a way that makes them appropriately trusted (or untrusted) subjects in a TCB state machine.

In the real world, however, this approach overlooks a problem called dynamic proliferation. Subjects and objects change state too quickly to cost-effectively maintain the TCB perimeter. The perimeter must expand and contract constantly if the enterprise is to function effectively within the business environment. With dynamic proliferation, each subject and each object must carry its own set of acceptable states, in effect forming a "mini-TCB" that must be carefully maintained. For example, a file could keep its own secure record of who can access it and what can be changed, with the record attached to the file itself and not stored in a separate database.

Currently, there are several initiatives targeted at addressing the disappearing perimeter. On the subject side, there are solutions such as federated identity management, identity and access management, and security information management. For objects, the individual repository/processing unit (the server) needs to function as an isolated TCB, which translates into the need for a secure platform.

5.3.3. Strategies for Implementation

Knowing the difficulties inherent in identifying and controlling the TCB, how can secure platforms be established from which to launch and manage connected enterprise services? The next few sections outline these steps.

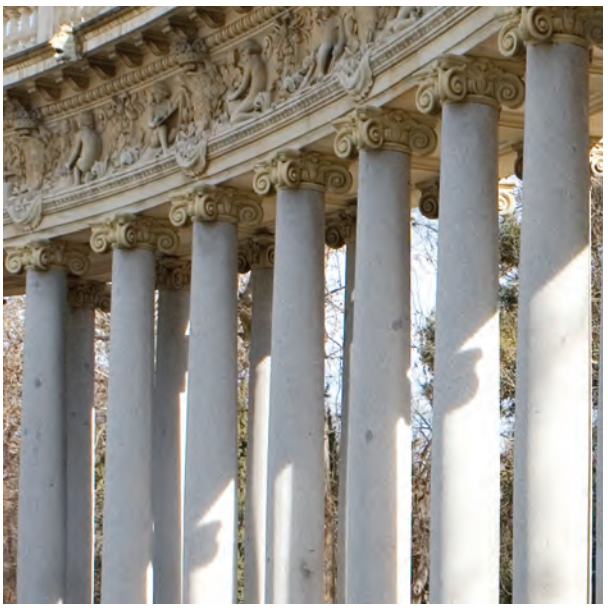

5.3.3.1. The Confidentiality, Integrity, and Availability (CIA) Triad

All security exists to ensure exactly three things, *confidentiality*, *integrity*, and *availability*:

- Confidentiality implies no unauthorized disclosure of information.
- Integrity implies no unauthorized modification or destruction of information.
- Availability implies that authorized users can access information when it is needed.

At its most basic level, platform security selects assets that can be confined to a single server and ensures that appropriate levels of confidentiality, integrity, and availability are guaranteed for the assets while they are on that server. Assets may be data or programs, CPU cycles or bits, subjects, or objects. Security analysis usually involves a large number of security goals, threats, threat agents, exposures, risks, and countermeasures. However, the analysis circles back to ensuring some combination of these three basic properties.

5.3.3.2. Identifying Vulnerabilities

Like most security analyses, the first step in planning a secure platform is to identify the realistic vulnerabilities relative to the value of the assets being protected. Using the CIA triad is particularly helpful in this case, because it helps to quickly sort subjects and objects, and it elicits a description of useful and not-so-useful state changes.

There are three key questions for identifying vulnerabilities:

- How can the confidentiality of information on this platform be compromised?
- How can the integrity of information on this platform be compromised?
- How can the availability of information on this platform be compromised?

Answering these three questions in detail requires the security architect or consultant to address a number of other questions as part of a standard risk analysis. For example, the questions above cannot be adequately answered without asking:

- What information is stored on this platform? That is, what are the assets?
- What does confidentiality mean in this situation?
- Who is authorized for what information, at what time, under what conditions?
- What does integrity really mean? Who is authorized to change data? What internal verification mechanisms are already in place that guarantee integrity or obviously identify data integrity issues?
- What does availability really mean? How many users should access how much data over what time span? How often does the data change, and how quickly must those changes be propagated?

There are also inductive vulnerability assessment techniques, which involve attacking the platform in question with various exploits to see how the confidentiality, integrity, or availability of the platform might be violated. However, these must be preceded by (at least) a rudimentary paper analysis. Without knowledge of what CIA means to the enterprise, it is difficult to gauge whether a given attempt is an attack or an acceptable access method.

5.3.3.3. Identifying Threats and Threat Agents

After assessing what CIA means for a given organization or enterprise, it is useful to evaluate the threats by separating CIA into a series of common security goals, for example:

- Maintaining privacy: Protecting from unlawful disclosure
- Maintaining secrecy: Protecting from industrial espionage
- Maintaining integrity: Keeping the data intact
- Maintaining access to service: Keeping the system up and running
- Limiting abuse: Defending against a malicious internal user
- Identifying problems: Overcoming stealth
- Assuring security: Locking out unauthorized users
- Maintaining security policy: Knowing what to do, when to do it, and how to do it

For each area, many different threats can occur that vary in type, format, and means of attack. Rather than cataloging the threats, each of the security goals is detailed as a means for easily recognizing potential threats.

Maintaining Privacy

Privacy of data (one aspect of confidentiality) must be maintained. Certain data must be kept strictly in confidence. The risk associated with the loss of privacy is known as unlawful disclosure. Each person and enterprise should have the opportunity to choose when and with whom data is shared. In many industries, such as telecommunications and medical services, regulatory requirements and disclosure laws provide stiff civil or criminal penalties for failure to maintain the privacy of data.

Unlawful disclosure usually occurs in one of four ways:

- An authorized party (responsible for maintaining privacy) reveals information through error, neglect, or malicious intent.
- An authorized party (responsible for maintaining privacy) accidentally or deliberately grants access to an unauthorized party.
- An unauthorized party monitors communications channels (for example, a telephone tap) to obtain information while it is transmitted between authorized parties.
- An unauthorized party obtains direct access to files or other information resources to collect information.

Maintaining Secrecy

Access to competitive data should be limited to a need-to-know basis. Data is usually classified into risk categories (For example, company-confidential or competition-sensitive), with access to a category tied to a title or position (role). Disclosure may be unintentional or malicious. Public disclosure of secret information can mean the loss of revenue and competitive edge.

When organizations become very large, it is usually impractical to explicitly identify each person who has access to competitive information. Instead, classifications (levels of secrecy) are used. These are typically connected to job description or position in the organizational chart. This kind of security is called multi-level security because there are need-to-know or safe-to-know strata that define who can know what. All such multi-level security measures are designed to reduce the probability that sensitive data will end up in the hands of a competitor or someone who will deliberately use it to damage the enterprise.

Maintaining Access to Service

Losses can be incurred because information or computing resources are not available. Deliberately preventing legitimate access is known as Denial of Service (DoS). A person or enterprise should not be prevented from using information because someone else maliciously disables the means to access that information. This also applies to information resources, such as computers, networks, and communications systems. Of all threats, DoS is the most insidious and the most difficult to prevent. A simple example is someone who ties up a competitor's telephone lines with bogus calls, preventing legitimate customers from being serviced.

One of the most common Internet DoS attacks, which can be performed by relatively unsophisticated attackers using tools available from certain web sites, involves constantly accessing (hitting) a site's homepage, causing some customers to time out without accessing the page. If a malicious organization employed enough agents, each using a web browser to repeatedly request a competitor's web pages, the target would be effectively closed down. Because the web pages must be available to everyone on the Internet to be effective, it is not possible to totally prevent this attack. However, effective security strategies can significantly reduce the impact and subsequent risk.

Maintaining Integrity
Deliberate corruption or destruction of data can deny access through:

• Outright destruction of files: Another variety of DoS
• Overt corruption of files: Data is obviously obliterated or garbled beyond usability
• Covert corruption of files: Data is altered in a way that is not immediately apparent to give false impressions
• Corruption of computer programs: Programs are modified to take unauthorized or destructive actions

A customer or enterprise should not be prevented from using information because someone else destroys it. This goal covers data that has been imperceptibly altered to produce bad decisions or false conclusions. It also addresses bogus programs that damage the system, including Trojan horses, viruses, and other forms of malicious code.

Limiting Abuse
Employees must not be allowed to betray trust by:

• Gaining unauthorized access to corporate data or computing resources
• Granting access to an unauthorized party
• Misusing corporate computing resources
• Corrupting or destroying computing resources

Privileged users must not be allowed to betray the trust granted to them by the organization. There are several ways that privileges can be abused.

Enterprise employees may gain unauthorized access to files or corporate information systems, accessing data for which they are not entitled. They may grant unauthorized data access to a third party, such as a competitor or foreign power. They may misuse corporate computing resources to perform essential services for a competitor; or they may simply corrupt, obliterate, or steal corporate resources, as in the case of a disgruntled employee.

As with DoS, there is no perfect defense. However, limiting employees' access to competitive data, confidential data, and resources not required for their job has a tremendous impact on mitigating this risk factor.

Identifying Problems
Identifying an attack is a cornerstone of layered security protection. Enterprises must know that a breach has occurred, identify the perpetrator and/or the means of attack (if possible), and quickly assess and control damage. Solid problem identification is the most significant step in damage control.

In spite of active security measures, there is always a probability (however small) that someone will penetrate the system. If the surveillance system is well-designed, however, the chances are high that a perpetrator will be caught or positively identified. In addition, the presence of visible surveillance often acts as a powerful deterrent to potential violators.

Even if the perpetrator is not identified and caught, enterprises must be able to assess and repair the damage as accurately as possible and repair the exploited vulnerability. This assessment is the most significant step in damage control. For example, if a corporation knows that its pricing strategies are compromised, it could change the data to confuse the perpetrator.

Assuring Security
A secure system is only part of the security solution. The system must also be configured, maintained, and operated properly. In addition, corporate procedures must support system security. Confused administrators and sloppy procedures are easy targets for attackers. To ensure that security and policy compliance is maintained, administrators must clearly understand the steps to take and the correct order. Confusion regarding the administration of a secure system often leads to inadvertent openings that a perpetrator can exploit. In addition, site security policy must ensure that hard-copy documents, media, and conversations do not reveal information being protected by the secure system. For example, positioning a computer screen to face an uncovered, first-floor window could easily defeat the purpose of all other security features.

Maintaining Security Policy
Security policy is important to the people and process part of the security equation (people, process and technology). Security policy is the set of rules and procedures for people in the organization to follow, and it also serves as a set of guidelines for process. Security policy spans how to handle information, how to conduct business transactions, what to do in the case of a security incident, and what happens when security policies are violated. To be effective, security policy maintenance must start with awareness and training, and it should continue with policy updates. All the while, documentation should also be maintained for legal and regulatory policies that require monitoring for compliance, enforcement, and investigation.

5.3.3.4. Assessing Risk and Choosing Countermeasures

Effective risk analysis for implementing a secure platform hinges heavily on the correct use of the CIA triad (discussed earlier in this chapter). It also relies on the careful and ongoing assessment of vulnerabilities, threats, threat agents, losses, exposures, and risks.

It is useful to define different types of risk-mitigation strategies that help to secure a computing platform or operating environment. In fact, there are some proven risk-mitigation strategies that help to meet the collection of platform security goals discussed previously. Furthermore, layering these strategies dramatically increases security and decreases risk. Risk-mitigation strategies include:

- Internet traffic filtering
- User authentication
- Data partitioning
- Integrity checking
- Use of least privilege
- User authorization
- System surveillance
- System alarms
- Simple security administration
- Clear site security policy, including compliance monitoring and enforcement
- Ongoing user training and awareness efforts

Internet Traffic Filtering

Stopping problematic traffic before it reaches a system averts subsequent problems and cleanup work. Filtering known bad traffic (such as virus attacks) and preventing inbound or outbound connections from or to known bad IP addresses are two examples of network traffic to stop at the outside edge of an infrastructure. A firewall is one technology that allows this type of filtering.

User Authentication

Many tools for guessing or cracking passwords are freely available. Given the low cost of powerful computers and the fact that many people choose easy to guess passwords, password cracking has become a very simple operation. To combat this, it is important to improve user authentication before granting access to resources. This can be accomplished using a combination of three authentication methods:

- Something the user knows: passwords that are improved to thwart cracking attempts. Password-hardening tools make users select passwords that are not comprised of common words and names.
- Something the user has: smart cards or physical token devices (such as a keychain security token) can respond to a challenge during login. Users login, enter the password, and the system challenges them to enter a valid ID number (or some other credential) from the smart card or token.
- Something the user is: biometrics refers to measurements of unique physical features of human beings including fingerprints, retinal scans, voice printing, and blood vessel printing.

There are different ways to authenticate that users are who they claim to be. Additionally, it is important to select the right level of authentication to meet security requirements and policies.

Data Partitioning

Access can be controlled by implementing a multi-level security system. Programs and users are given a clearance, and files and data are given a label. If the label does not match the clearance, access is denied. Hierarchical access can also be defined in such schemes. Multi-level security systems partition data into compartments, for example, inside (intranet) and outside (Internet). Programs running on the outside cannot access files on the inside, and vice versa. Attackers coming in from the Internet should not be able to reach into the inside compartment to access data files, run programs, or download/upload files. If programs in two or more compartments must share data, hierarchical access may be necessary. In the example given, a system compartment may be required to store configuration files and other files needed by all system programs and applications.

Integrity Checking

Security features are typically complemented with integrity checking, for example:

- Files, directories, and system tools carry security attributes.
- A master list of security attributes is maintained in a safe location.
- Files and directories are periodically checked against the list.
- If discrepancies are found, they must be explained and then fixed.

In an integrity-checking system, the system knows which security attributes (for example, owner, data-partitioning compartment, read access, checksum, or signature) should be assigned to key files. An administrator runs the program periodically to check the state of the file system. Any errors are immediately flagged, and the administrator can reset file attributes, restore from known good copies (if available), or disable the system until an investigation can take place.

Use of Least Privilege

In nearly every OS, programs use such OS services as terminal input/output (I/O) for portability. Normally, every program has access to every service. This unlimited access presents an opening for rogue programs and hackers. To restrict behavior, each service is protected with a privilege. If a requesting user or system has insufficient privilege(s) for access to a particular service, then that service should not be accessible. Using only the needed privilege for the shortest possible time is known as least privilege.

Every program that runs on a system must perform certain basic tasks. Because programmers do not want to recreate all of these basic operations, such as accessing files or controlling a display, the system provides a set of system services (sometimes called system calls) to handle them. Access to system services is typically unrestricted - meaning that an attacker's malicious program could easily use system services to bypass security measures.

To overcome this weakness, system services can be divided into classes, with a specific privilege assigned to each class. All programs that expect to perform a basic operation such as I/O must present the appropriate privilege for each operation and then relinquish the privilege when it is no longer needed. For example, if a program needs to store a file, it should request write permission to the destination file system, perform the store function, and then relinquish write permission. Therefore, the program only possesses the ability to write to that file system for the time required. If the process was subsequently compromised, it would not have the ability to write.

Privileges significantly reduce the level of risk from malicious programs and Trojan horses. Because privileges must also accommodate off-the-shelf or legacy applications, there must be a special category of privileges for executable programs. However, these privileges should only be assigned by a properly authorized user within the enterprise.

User Authorization

Many systems have a super user or root account that is all-powerful and to which all administrators have full access. Unlimited access provides an opportunity for attack. Superuser access must be divided into discrete authorizations. Every administrative job or role carries a subset of these authorizations. If defined correctly, these roles provide a balance of power.

The most common means of penetrating an HP-UX-based system is to obtain the super user account password. If the various capabilities normally assigned to the root account are divided into discrete authorizations, each of which permit access to a very limited subset of capabilities, the super user account can be disabled or restricted to only a few highly trusted individuals. Each discrete authorization allows access to a very small set of system features. Authorizations can be grouped into roles and distributed to specific individuals, so that users or systems have only the authorizations required to perform their specific functions. For example, night operators do not usually need super user access to the system, but they do require certain access privileges beyond that of a normal user.

System Surveillance

There is no substitute for monitoring and surveillance. Effective monitoring and surveillance should:

- Execute at a low level, within or near the OS
- Record system events with timestamps and user IDs (auditing)
- Avoid degrading performance by allowing tuning and customization
- Collect and present information in real time, if possible

Because most security breaches involve stealth, the system should notify the appropriate administrators and/or security personnel when security has been breached. This allows personnel to quickly assess and limit the damage. To provide this information, the system should implement an auditing system to log suspicious activities. Since any system activity, maliciously used, could be considered suspicious, these activities must usually cover the full range of processes at the OS level. Monitoring at the system level minimizes the chances of disguising suspicious activity.

Because OS-level activities represent a very large volume of audit data, there must be a mechanism to tune resource usage, such as disk and CPU time. In addition, audit trails must often be preserved as evidence. Records management utilities must be available to allow audit data to be offloaded to and restored from removable media on a session-by-session basis. Finally, there must be post-processing tools for analyzing audit data and producing reports. Various kinds of filtering tools are needed to help focus the search for suspicious behavior.

System Alarms
Because auditing is passive, active surveillance is also needed. Effective active surveillance should:

- Execute at a low level, near the OS
- Monitor audit events or key system activities
- Filter data to allow targeting of specific events or times
- Provide real-time notification
- Activate automated defense measures
- Prioritize responses (if and as appropriate)
- Offer a high level of user configuration

Audit data must be examined carefully and the information is relatively detailed. To ease the data analysis burden and provide a real-time intrusion detection capability, systems can implement an alarm capability. Alarms can use the same set of system events recognized by the auditing feature, or they can use pseudo-events that address common penetration points.

Alarms can provide alerts for:

- Specific events that occur all the time
- Events that happen at an unexpected time of day (for example, a login at 3:00 in the morning)
- Events that happen too often (for example, five consecutive failed logins)

Alarms can also be implemented to select only certain conditions or patterns on which to trigger. When an alarm is activated, the system may do anything from simply logging the alarm to paging an operator and shutting down the system. The actions depend on the system configuration and security policy in effect at the site.

Simple Security Administration
Poor administration can nullify even the most effective security. Important attributes of security administration include:

- Security features must be simple to administer.
- Administration must be similar in format to normal operations.
- Steps must be clear and well established.
- Training of appropriate personnel should be thorough and ongoing.
- Updates to protection tools (such as patches and threat signatures) should be tested and applied in conformance with clear site security policies.

Key security features must be controllable from a native system interface, in a format consistent with the normal functionality of the system. For example, a menu-driven system should have menu-driven security controls. These controls should be divided into categories corresponding to security roles defined at the site (for example, operator, night operator, and system manager). Because maintaining a secure system can be a complex process, online help and documentation is usually essential. Each step that must be taken to ensure security must be represented in order and in a format that is easily accessible to administrators.

Clear Site Security Policies

HP has well-established processes for defining and developing security policy. It is sufficient to note some areas that are often under-addressed when dealing with platform security:

- Physical handling of media and hardcopy
- Physical access rules and procedures
- Platform security configuration control and update policies
- Handling of suspected or known penetration attempts (incident response)
- Training
- Policy compliance monitoring and enforcement

The key is to include the platform and its unique security management requirements in the overall security policy analysis.

5.3.4. Device Security in Practice

Virtually all systems in the infrastructure have information that is important to users, and the task of reconstructing data and applications is never pleasant. A lost or stolen device, damage caused by virus infections and the need to get a user or an application back to work as quickly as possible, can all place a significant amount of stress on those needing to resolve the situation. Proactively minimizing the potential of these types of issues can make a tremendous difference in resources expended for safeguarding devices and the information assets they must protect.

The first step in constructing such a plan is to understand the organization's specific risk environment and supporting policies and will often include a combination of asset control policies and procedures, which include topic areas such as physical security, device backup or synchronization, standardization for mobile platforms, authentication, storage, and encryption. Having assessed these areas, the next logical step will be to identify individual solutions that will efficiently and effectively assist in managing the environment form a holistic perspective. HP offers a solution with HP ProtectTools and HP Secure Advantage for HP client, server and storage system portfolios.

Protecting data at rest on devices across the infrastructure is the first step of a comprehensive security architecture.

5.3.5. Protecting Data at rest

Protecting data at rest requires protecting against threat of unauthorized access to sensitive information stored locally on the device itself. While this has traditionally been addressed with procedural measures for device access and administration, the increased complexity of data centers and the growing number of mobile and remote client systems is making it essential to build access control and data protection mechanisms into the infrastructure itself. Data protection solutions should secure the device with tight access control mechanisms, whether at the OS level, with data encryption, or better yet, using supporting hardware security features. Recurring, high-profile incidents spanning industry and government agencies reinforce the dangers inherent in allowing sensitive data to reside on user-controlled systems. To protect sensitive data, organizations must develop, teach, and motivate individuals to strictly adhere to policies for transferring sensitive data from central repositories to individual systems, as well as deploy infrastructure technologies that provide strong controls and protections to data at rest on client or server systems. HP security solutions provide means of achieving this by deploying encryption solutions across client and server systems, with support from embedded hardware security solutions that are designed to the Trusted Computing Group specifications.

Strong user authentication can help access control go beyond strong passwords, and move to hardware-based tokens such as smart cards. Requiring multi-factor authentication, such as biometrics combined with passwords is an important means to increase device protection. Power-on authentication, such as DriveLock technology, provides additional protection. HP's client portfolio integrates strong user authentication solutions with power on authentication, disk access control mechanisms (such as DriveLock), and full volume encryption solutions, with support from embedded hardware security such as a TPM. These integrations provide stronger user authentication, and they aim to tie local data on the hard drive to a particular client device for enhanced data protection.

5.3.6. Mobile Device Security

Mobility has extended the device spectrum from traditional desktops and servers to notebooks, handhelds, phones, and a wide range of specialized appliances. These devices are vulnerable to a new set of security issues, including susceptibility to loss and theft, increased use outside company premises, and less processing power to ward off threats. Generally, mobile device security falls into three areas:

- Securing local data from unauthorized access
- Safeguarding the device from malicious threats or data loss
- Protecting connectivity between the device and the applications residing on corporate servers

5.3.6.1. Securing Local Data from Unauthorized Access

The most common concern relating to mobile device security is the threat of unauthorized access to sensitive information stored locally on the device itself. A fundamental distinction of mobile devices is that they can be used offsite in public (even adversarial) environments. Requirements for any solution should include securing the device with tight access control mechanisms, whether at the OS level, with data encryption, or with hardware features. Recurring, high-profile incidents spanning industry and government agencies reinforce the dangers inherent in allowing sensitive data to reside on user-controlled systems. To protect sensitive data, organizations must develop, teach, and strictly enforce policies for transferring sensitive data from central repositories to individual systems. Protection of data at rest will be achieved using mechanisms such as those described earlier in this section.

5.3.6.2. Safeguarding the Device

Like other systems, safeguarding applications and data is an issue for mobile devices. Virtually all systems have information that is important to users, and reconstructing the configuration of the applications and platform is never a pleasant task. In addition, there is also the need to cope with a lost device or a virus infection and get the user back to work as quickly as possible. Therefore, safeguarding devices calls for a multifaceted approach involving some combination of asset control policies and procedures (for example, what to do if a PDA is lost), physical security, device backup or synchronization, mobile platform standardization, strong authentication, and storage encryption. Of course, the first step in risk management is always to understand the organization's specific risk environment.

5.3.6.3. Protecting Connectivity

Most mobile users want to connect to their enterprise networks and access applications. These connections and actions need to be protected beyond the scope of the mobile device. Users need to establish a secure connection over a typically wireless transmission medium. In some cases, the device can connect directly to the private network, but oftentimes its path traverses some type of public network.

Popular public packet data networks include Wireless LANs (WLANs), General Packet Radio Service (GPRS), Cellular Digital Packet Data (CDPD), or even circuit-switched connections dialed in to an ISP. If IP networking is supported, the user may establish an IPsec or SSL VPN connection to the corporate infrastructure. For detailed information about wireless security technologies, including Wireless Personal Area Networks (WPANs), WLANs, and Wireless Wide Area Networks (WWANs), refer to the Trusted Infrastructure Network Security section of this handbook.

5.3.6.4. Centralized Policy Management

Covering all three areas of mobile security is an additional dimension that is particularly important for large organizations. Centralized policy management is possibly the most critical challenge that enterprises have in deploying mobile devices. A sound authentication and encryption solution should be enough to thwart all but the most resourceful of hackers. However, in an enterprise the concerns run deeper than the presence of a few software components. It is important to ensure that all of the tools and settings are configured correctly.

This requirement highlights the difference between consumer and enterprise concerns around mobile devices security. While individual customers have full jurisdiction over their own property, enterprises do not necessarily trust their users to configure and run their devices in a secure fashion, particularly when sensitive information is stored on the device.

Users may find strong passwords difficult to remember and periodic re-authentication to be annoying. If it is a user's personal device with only personal information nobody can blame the user if he or she chooses to lower his or her security settings. However, when the device contains critical enterprise data, security policies mandate that much stricter settings must be enforced.

One aspect of centralized policy management is the initial deployment of any security software. It also includes the ability of the enterprise to centrally deploy new applications and new versions of existing security products.

But simply setting up the device is not sufficient. It is also of critical importance to the enterprise to be able to ensure compliance of any devices that that might have sensitive data on them. This means that the enterprise needs to be able to determine if a terminal has been adequately security before allowing the user to synchronize any sensitive data.




5.4. HP Host Security Products and Solutions

HP offers a complete range of host security products and solutions to help address the threats and mitigate the risks discussed in this section. The following sections provide an overview of individual products and solutions.

5.4.1. HP ProtectTools for Client Security

As computers become more mobile and better connected, threats to data security are increasing in magnitude as well as complexity. Organizations in which data security directly impacts business health are becoming increasingly concerned about this problem. Client devices including notebooks, desktops and workstations tend to be the front line of access to an organization's information assets. As such, client device security becomes a key mechanism to securing the IT infrastructure. Security requirements at the client device level can range from strengthening user authentication, to hardening the client device (at the hardware, OS, or application level), to protecting data as it resides on the device.

Client device security is of strategic importance to HP, as it is to an increasing majority of business and IT managers. As such, HP offers the comprehensive HP ProtectTools client device security solution set. HP ProtectTools originated with an HP developed smart card security solution for client PCs. The application is now part of HP's business notebook, desktop, and workstation smart card solutions. As the HP security portfolio has grown, the HP ProtectTools name has also grown to represent a broad security solution set that encompasses software, hardware, and services. HP's Personal Systems Group (PSG) and HP Services (HPS) deliver security solutions designed to address security challenges at all levels of client devices.

5.4.1.1. HP ProtectTools Security Manager

Taking a holistic approach to security, HP designed the HP ProtectTools Security Manager to bring many technology areas together in a way that not only protects client devices but also prevents them from becoming points of vulnerability that threaten the entire IT infrastructure. The HP ProtectTools Security Manager is at the heart of the HP ProtectTools security offering for HP notebook, desktop, and workstation PCs. (See **www.hp.com/go/security** for product information.) This single-client console application unifies the security capabilities of HP client PCs under a common architecture and single user interface. A range of features build on underlying hardware security building blocks, such as biometrics, smart card technology, and embedded security chips (TPMs) designed in accordance with the TCG standard. Collectively, these features address business needs for better protection against unauthorized PC access and stronger protection for sensitive data.

Most importantly, HP ProtectTools hardware security mechanisms provide the enhanced benefit of not relying solely on OS and application security vulnerabilities that are known targets to most off-the-shelf hacking tools. HP ProtectTools Security Manager embodies an extensible framework designed to enhance security software functionality through add-on modules including:

Embedded Security for HP ProtectTools
Embedded Security enables strong hardware-based protection of data and digital signatures and reliable hardware-based device authentication.

Java Card Security for HP ProtectTools
Java Card Security enables stronger user authentication, using two-factor authentication and HP-patented pre-boot authentication technology.

BIOS Configuration Security for HP ProtectTools
BIOS Configuration Security provides easy access to features such as power-on user and administrator password management and easy configuration of pre-boot authentication features including smart card, power-on password, and TPM-embedded security chip.

Credential Manager for HP ProtectTools
Credential Manager provides flexible multifactor authentication that combines a wide array of devices including biometrics, smart cards, and USB tokens, and provides a completely automated single sign-on with automatic field detection, registration, and credential entry.

Device Access Manager for HP ProtectTools
Device Access Manager creates policies that control which users or user types get access to which devices or device types. With Device Access Manager, policies can specify what types of devices are allowed or disallowed, depending on the user, such as allowing keyboard and mouse but disallowing USB storage devices for users without administrative rights. An enterprise version of Device Access Manager is also available for configuring and deploying the same policies remotely.

Drive Encryption for HP ProtectTools
Drive Encryption encodes every bit of information on your hard drive volume so that it becomes unreadable to an unauthorized person. This feature helps ensure that sensitive information cannot be accessed if the notebook, desktop, workstation or hard drive is lost or stolen and also ensures critical personal and business data stored on the hard drive is safer without user intervention. Drive Encryption is a standard feature on select HP business notebooks, desktops and workstations.

5.4.1.2. HP ProtectTools for Microsoft Products

The HP Microsoft ProtectTools suite adds to the security functionality provided by standard Microsoft products. They have been developed in partnership with Microsoft to ensure they integrate seamlessly with standard Microsoft products in order to meet the needs of security-conscious organizations. These HP ProtectTools Secure Commercial Off The Shelf (SCOTS) products have the following functionality:

HP ProtectTools Authentication Services
This product provides a number of features that enhance the standard Microsoft authentication process. The central feature is enhanced password management, achieved by implementing a CESG-approved password hashing and password generation system. Each organization is provided with special CESG seed values to ensure each organization's system is unique. Where government algorithms are not applicable, alternative commercial algorithms are used.

The product also manages change of administration passwords, provides last successful and unsuccessful login information, and can be configured for multiple login denial and timed auto-logout. The use of a unique password hashing mechanism for systems prevents access from unauthorized systems even when a valid username and password are used.

HP ProtectTools E-mail Release Manager
The inappropriate release of a sensitive e-mail is a constant threat to any organization. E-mail is the easiest way of sending information around the world via the internet, and once the e-mail is sent, it has gone. HP ProtectTools E-mail Release Manager integrates with Microsoft Outlook to help mitigate such threats. For example, every e-mail created could be made to automatically carry the label of "Company Confidential". Any e-mail with this label could then be restricted to a distribution of employees with the relevant authority. Other labels could be configured for specialist teams, the general public, related organizations or suppliers, senior management only or "Project X only" groups, respectively. Any, or all, of these could have the e-mail electronically signed, encrypted, and audited, and additional user input can be mandated to confirm the authority to send that type of e-mail. The HP ProtectTools E-mail Release Manager Outlook Web Access extension provides the same functionality but in a web-based environment.

HP ProtectTools Device Access Manager
Device Access Manager for HP ProtectTools (DAM) is a plug-in module into HP ProtectTools Security Manager (which is discussed in section 5.4.1.1.). HP Security Manager is a single-client console application that unifies the security capabilities of HP client PCs under a common architecture and single user interface. HP Device Access Manager can be found within the HP ProtectTools Security Manager option in the Control Panel. It provides a Simple Configuration view, which offers the most common access scenarios, and an advanced Device Class Configuration view allowing more complex access scenarios to be specified. Access to the configuration can be controlled and is restricted to Administrators or delegated to power users via the User Access Settings configuration view, and only authorized users can access the Security Manager DAM utility and modify the configuration.

The HP Device Access Manager for ProtectTools is designed for standalone use, whereby all of the configuration settings are stored locally in the registry. The related management product, Enterprise Device Access Manager (EDAM), stores configuration in the Active Directory.

HP ProtectTools Enterprise Device Access Manager

Managing and controlling the import and export of data onto or from systems is essential. Today's systems come with floppy drives, CD/DVD read and write drives, compact flash cards and a combination of USB, serial and parallel ports. In many cases these devices are essential for carrying out day-to-day business, but they can present security threats as they provide means of exporting confidential data or of importing malicious code. One option is to have devices removed from the system altogether, but this is expensive and results in non-standard system support and maintenance. HP ProtectTools Enterprise Device Access Manager (EDAM) controls access to classes of devices or individual types of devices based on permissions granted to a computer, to a user, or to a group. HP ProtectTools Enterprise Device Access Manager utilizes the Microsoft Windows Active Directory to store and propagate device access permissions throughout the Windows domain. It can even control the type of devices that are allowed to connect to a particular port. For example a user may be permitted to connect a particular printer, mouse and keyboard via a USB port while excluding any mass storage device.

HP ProtectTools Role-based Access

In today's busy working environments many people undertake numerous different jobs or roles. These roles need different kinds of access to different systems or applications. HP ProtectTools Role-based Access provides the opportunity for a standard desktop across a whole organization while providing secure terminal server access from any workstation to a user's applications and data for their particular role or roles.

HP ProtectTools Windows Mobile

Most organizations see mobile commuting as the next big opportunity for achieving cost reductions while increasing business efficiency. Security has been the major concern preventing the take-up of this technology. HP has the capability to secure remote connections and protect the data held on mobile devices such as notebooks and PDAs. HP ProtectTools Windows Mobile ensures that proper authentication is undertaken, that all data is deleted if the system is lost or stolen and the password is incorrectly entered a predetermined number of times, and ensures that PDAs can only link with known and authorized PCs.

HP ProtectTools Application Manager

A major concern for all organizations is the threat from malicious code that is introduced into computer systems known as malware. Such malware can be introduced inadvertently by honest users or deliberately by malicious users. This can happen through internal network connections, internet access or a host of different peripherals such as CD/DVDs, floppy disks, memory sticks, etc. HP ProtectTools Application Manager acts against this threat by controlling the execution of code. In the default configuration, only code introduced by the system administrator, or belonging to the system account, will be allowed to execute. This configuration can be extended to require signatures on executable files which are validated as the file is loaded and before they can be run. The fundamental approach of checking ownership and the signatures of the files, before allowing them to execute, protects against malware from all sources. This same tool can also be used to control access to executable code already loaded which requires controlled access such as system configuration programs.

With the HP ProtectTools for Microsoft Products suite, HP addresses the security needs of a broader range of market segments. Customers benefit from HP security innovations that have proven reliability - the result of an exhaustive validation process in demanding customer environments.

For more information, see the HP ProtectTools overview and white papers at **http://h20219.www2.hp.com/services/cache/45782-0-0-225-121.aspx**. Free software evaluation downloads are available at **www.software.hp.com**.

5.4.2. HP NetTop

HP NetTop is a highly secure and layered architecture of Security-Enhanced Linux (SELinux), the VMware Workstation, and customized security policies. HP NetTop is an information assurance solution that enables connectivity to multiple network domains of differing sensitivities from a single system, while maintaining data and domain separation through secure virtual machine air gaps.

HP NetTop is backed by the HP Technology Solutions Group to provide assessment, planning, policy definition, rollout, and support tailored to an organization. HP NetTop provides strong compartments that meet many government and financial industry requirements. Originally developed by the National Security Agency (NSA), HP NetTop is now offered by HP as a full-service solution to public and private enterprises.

5.4.2.2. HP NetTop Solutions

HP NetTop solutions exist for both public and private organizations. HP provides:

• Health care organizations with HIPAA compliance by maintaining patient records in isolated domains while allowing access to those who need it
• Financial institutions with customer record and financial data security
• U.S. Defense and intelligence agencies with Director of Central Intelligence Directive (DCID) 6/3 Protection Level 4 (PL4)-compliant, low-cost security domain separation and access to multiple coalition networks

A complete security solution - from initial assessment through rollout, training, and post-deployment support - ensures that HP NetTop works now and in the future. With HP's unified desktop and delivery, HP NetTop adapts enterprise computing to current and emerging risks. For more information about HP NetTop, visit www.hp.com/go/nettop.

5.4.3. Host Operating System Security

5.4.3.1. HP-UX

For the past 20 years, HP has delivered one of the most trustworthy and secure UNIX operating systems. Designed for protection against both external and internal threats, the HP-UX 11i operating system has a well-integrated set of security features aimed at proactively mitigating risk and helping reduce compliance cost.

For enterprise customers who must respond to constantly changing business needs, security solutions provided with the HP-UX 11i operating system simplify the deployment of layered security features while providing the extra assurance of in-depth protection. HP-UX 11i security solutions are included as part of the base HP-UX operating environments.

5.4.3.1.1. Platform Security Fundamentals

The most basic goal of operating system security is to preserve the integrity of the system in the face of attack. The HP-UX 11i operating system includes a number of features that assist the administrator in locking down the platform:

• **HP-UX Bastille** provides a graphical interface that guides an administrator in tasks that harden the system against attack, including locking down system ports, files and other components.
• **Host IDS** uses kernel-level system audit information to continuously monitor many systems for attacks, generating alerts and, as an option, also responding in real-time.
• **IPFilter** provides system firewall capabilities, including stateful connection filtering to limit the "attack surface" of the platform, and connection throttling to limit the effectiveness of DoS attacks.
• **Install-time Security** eases default lock - downs by offering a menu of security profiles that may be applied as part of the OS installation process.
• **Security Patch Check** helps ensure security patch currency by periodically connecting to HP and recommending the latest security-relevant patches.
• **Execute-protected Stack** prevents common types of buffer overflow attacks, which are a leading contributor to platform compromise.

5.4.3.1.2. Security Containment

HP Security Containment for HP-UX 11i is a suite of security technologies designed to dramatically reduce the likelihood of system compromise. HP incorporates these enhanced security features into the mainstream HP-UX 11i operating environment to help businesses combat increasingly complex threats. Without requiring modification to applications, HP Security Containment isolates compromised applications, which are denied unauthorized access to other applications or files on the system. HP-UX 11i Security Containment comprises three core technologies that together provide a highly secure operating environment:

• **Compartments** provide isolation and restrict access to application and system resources outside of the compartment to prevent catastrophic damage should a compartment be penetrated. HP-UX Security Containment accomplishes this by controlling the flow of information between processes in different compartments. For example, outside compartments can accept and process customer-facing data, and then transfer it securely, by rule, to inside compartments for non-public access and processing.

• **Fine-grained Privileges** grant only the privileges needed for a task, and optionally, only for the time needed to perform the task. Applications that are privilege-aware are able to elevate their privilege level during the operation and lower it after completion of the operation.
• **Role-based Access Control** provides a mechanism to allow non-root users to perform administrative tasks, effectively splitting the power of root into a manageable set of roles. An out-of-the-box configuration supports many common HP-UX 11i commands.

5.4.3.1.3. Mission-critical Virtualization

The HP Virtual Server Environment (VSE) enables companies to take advantage of consolidation and virtualization techniques to improve server utilization while reducing operating system management costs and increasing security levels. VSE provides a mechanism for consolidating applications within a single operating system image. By combining the benefits of HP Process Resource Manager for resource entitlement and HP Security Containment for security isolation, Secure Resource Partitions provide a secure solution for lightweight virtualization within VSE infrastructures.

5.4.3.1.4. Identity Management and Accountability

HP-UX 11i v2 provides a number of built-in features designed to support the implementation of identity management architectures to provide manageable access control policies.

• **Standard Mode Security Enhancements** offer granular account and password policies on a system-wide or per-user basis, including the ability to generate detailed system audits for user accountability.
• **HP-UX LDAP-UX client services** simplify identity management by allowing system authentication and naming services to leverage a new or existing LDAP directory.
• **Kerberos** server and clients offer enterprise-class SSO services as well as enhanced interoperability with Windows ADS.
• **HP-UX AAA server** (RADIUS) authenticates network devices and controls access.
• **Red Hat Directory Server for HP-UX** provides an industry-standard, centralized directory service to store digital identity information.

5.4.3.1.5. Common Criteria Certification

The HP-UX 11i v2 operating system running on HP 9000 or HP Integrity platforms has been successfully evaluated against the requirements for the EAL4 Common Criteria (ISO 15408) Assurance Level, augmented by ALC_FLR.3 (flaw remediation), using the Controlled Access (CAPP) and Role-based Access Control (RBAC) Protection Profiles. EAL4+ is sometimes used as the abbreviated form for additional assurances. Details of the evaluation and evaluated configuration are available at: www.commoncriteriaportal.org/public/files/epfiles/CRP225.pdf and www.commoncriteriaportal.org/public/files/epfiles/hp-ux11iv2.pdf.

5.4.3.1.6. HP-UX 11i v3 Enhancements Strengthen Security and Streamline Compliance

HP is committed to a long-term roadmap for the HP-UX 11i operating system that encompasses continued enhancements to its built-in security features. The latest release of the operating system, HP-UX 11i v3, introduces a suite of new features that proactively mitigate risk and reduce the cost of compliance:

• **Encrypted Volume and File System** transparently protects data at rest against unauthorized disclosure if the data is lost or stolen, and may also provide safe harbor, avoiding the need for breach disclosure required by some state breach disclosure laws.
• **Trusted Computing Services** provides software support for Trusted Platform Module (TPM) embedded security hardware that is available on select HP Integrity servers for enhanced key protection and EVFS auto-boot support.
• **HP Protected Systems** offers an automated mechanism to configure and deploy more secure systems by leveraging the built in protection of HP-UX 11i servers, reducing the time and level of security knowledge required by IT personnel when configuring such mechanisms as Security Containment, which isolates processes and resources.

• **HP-UX Bastille** with drift reporting checks the consistency of a system's hardening configuration with previously applied hardening policy to avoid risk of system changes. This data reduces system exposure to malware, simplifies compliance maintenance, and provides visibility into undone hardening to allow planned response without risk of unexpected system breakage.
• **HP-UX AAA Server** offers more flexible integration with enterprise databases in combination with centralized, RADIUS-based user authentication and network access logging to simplify auditing and compliance.

The HP-UX 11i operating environment provides a comprehensive array of features that automate security processes, mitigating risks and lowering the cost of compliance. HP rigorously designs, engineers and tests these features through targeted development as well as collaborative projects with open-source and third party partners. Fully integrated within the HP-UX 11i operating system, this continually evolving suite of security enhancements is available at no extra cost to HP-UX customers.

5.4.3.2. Microsoft Windows and Server Applications

HP Services provides a number of offerings and security services (including design, configuration and hardening services) around Microsoft OS security in a client role and in a number of server roles. The Microsoft Windows Server OS includes a number of server applications such as Domain controller, Directory server (Active Directory (AD) or Active Directory Application Mode (ADAM)), Dynamic Host Configuration Protocol (DHCP) server, Domain Name System (DNS) server, Windows Internet Naming Service (WINS) server, File server, or Print server, Internet Information Server (IIS) (Microsoft's application server), Internet Authentication Service (IAS) server (Microsoft RADIUS server), Certificate server (for Public Key Infrastructure services), Network Access Protection (NAP) server (Microsoft's Network Admission Control (NAC) solution).

Other specific server roles (not bundled with the Windows Server OS and part of dedicated Microsoft software offerings) include Exchange Server (Microsoft's messaging server), Office Communications Server (OCS) (Microsoft's real-time collaboration server), SharePoint Portal Server (Microsoft's web portal server), SQL Server (Microsoft's database server), Identity Lifecycle Manager (ILM) (Microsoft's identity management and provisioning server), System Center (Microsoft's management server), and BizTalk Server (Microsoft's business integration and process management server).

HP also offers exclusive Microsoft Windows NT, Microsoft Windows 2000, Microsoft Windows XP, Windows Vista and Microsoft Windows Server 2003 security solutions. These solutions can for example replace the password hashing algorithms supplied in Microsoft Windows with customer-specific algorithms that make brute force or dictionary password hacking much more difficult. See also the "HP ProtectTools for Microsoft Products" section earlier in this chapter (section 5.4.1.2.) or the following URL for more information on the HP ProtectTools security solutions for Microsoft platforms and applications:
http://h20219.www2.hp.com/services/cache/45782-0-0-225-121.aspx.

Microsoft has also bundled an important set of host protection security features in its latest client (Windows Vista) and server operating systems (Windows Server 2008). These features are summarized in table 4-1 and explained below.

Table 4-1
Windows Vista malware protection features

Malware Prevention Features	Malware Isolation Features	Malware Remediation Features
• Security Development Lifecycle • Kernel-mode driver signing (64-bit only) • Patchguard (64-bit only) • Data Execution Protection (DEP) • Address Space Layout Randomization (ASLR) • Windows Defender • Automatic Updates • Windows Firewall • User Account Control (UAC) • Built-in Administrator Account Protection	• Service Hardening • Network Access Protection (NAP) Client • Windows Firewall • Internet Explorer Protected Mode • User Account Control (UAC) • Windows Defender • Windows Integrity Controls (WICs)	• Windows Defender • Security Center • Malicious Software Removal Tool (MSRT)

Fundamental Protection in Windows OSs

In Windows Vista and Windows Server 2008 Microsoft pioneers a couple of very fundamental malware protection measures that are not only related to new features included in these OSs, but also to the way the OSs were developed and engineered.

Windows Vista and Windows Server 2008 are the first Microsoft OSs that were developed following the Microsoft Security Development Lifecycle (SDL) methodology. SDL's primary goal is to improve the overall security quality of Microsoft software and make it more resistant to withstand malware attacks. SDL defines a formal and repeatable methodology that developers can leverage before releasing their code. Among the key elements of SDL are techniques for attack surface reduction analysis and measurement, and guidance for least privilege and security testing. More information on SDL can be found at this URL: **http://msdn.microsoft.com/security/default.aspx?pull=/library/en-us/dnsecure/html/sdl.asp**.

A Vista and Windows Server 2008 malware protection feature that is linked to software development, and more particularly to the development of drivers, is driver signing. Even though earlier Windows versions have an unsigned driver detection mechanism that can warn the user when he/she is about to install an unsigned driver, there were no driver signature checks on the kernel level. The new kernel-level driver signing and checking mechanism can (indirectly) better protect Vista from crashes or vulnerabilities that occur when malware installs or loads malicious drivers into kernel mode. More importantly, Windows Vista and Windows Server 2008 kernels require Windows Hardware Quality Labs (WHQL)-signed drivers: this means that the drivers are only signed after they passed a set of predefined tests that are run by Microsoft or one of its affiliates. More information on Vista driver signing can also be found at this URL: **www.microsoft.com/whdc/winlogo/drvsign/drvsign.mspx**.

At the kernel level, another Vista and Windows Server 2008 malware protection feature is called Patchguard. This feature is also referred to as kernel patch protection. Patchguard can prevent kernel-mode drivers from extending or replacing OS kernel services, and prohibit software from performing unsupported patches in the kernel. With Patchguard Microsoft is specifically targeting rootkits. Rootkits are software tools that try to conceal running processes, files or system data from the operating system in order to avoid detection.

Both driver signing and kernel patch protection are only implemented in the 64-bit versions of Vista and Windows Server 2008. Microsoft found that implementing these features in the 32-bit Vista and Windows Server 2008 versions was too difficult - not to say impossible. One reason for this decision was that most legacy 32-bit Windows drivers are not identified using a digital signature. Implementing stricter control over these modifications in the 32-bit Vista and Windows Server 2008 versions could have created major compatibility and performance issues for these legacy applications.

Data Execution Protection (DEP) is a memory protection feature that protects Windows systems against buffer overflow attacks - a technique often used by malware to compromise a computer system. During a buffer overflow attack malware tries to insert and execute code from non-executable memory locations. DEP allows Windows to mark certain memory locations as non-executable (NX). These NX memory locations can only contain data and the processor and the OS will prevent applications or services from loading executable code in them. DEP is not only supported in Vista and Windows Server 2008 but also in Windows Server 2003 Service Pack 1 (SP1), R2 and Windows XP Service Pack 2. DEP leverages a processor feature that AMD refers to as the no-execute page-protection (NX) feature and that Intel refers to as the Execute Disable Bit (XD) feature. At the time of writing AMD only supported NX on its 64-bit processors. Intel only supported XC on the Itanium and EM64T 64-bit processors and a small number of 32-bit Prescott processors.

To check whether a system supports this hardware-enforced DEP follow the procedure outlined in the following Microsoft Knowledge Base (KB) article: **http://support.microsoft.com/kb/912923**. Microsoft has also added a workaround - referred to as software-enforced DEP - that allows the Vista, Windows Server 2008, Windows Server 2003 SP1 and R2, and Windows XP SP2 operating systems to provide DEP on 32-bit processor systems. In this workaround the processor-level NX-or XD-bit is provided by set of cookies (or canaries as Microsoft refers to them) that the OS automatically adds to data objects stored in the OS heap and stack. See the following Microsoft KB article for more info on DEP and how to configure it: **http://support.microsoft.com/kb/875352/en-us**.

Ensuring Isolation in Windows OSs

Windows Vista and Windows Server 2008 include important new features to isolate the OS, services and data thereby making the platform more resilient to malware attacks. These features are the enhanced Windows Firewall, service hardening, inclusion of the Network Access Protection (NAP) client and Windows Integrity Controls (WICs).

A properly configured personal firewall is an important first line of defense for isolating an OS and preventing malware from infecting computers and spreading out across the network. Windows Vista's and Windows Server 2008's personal firewall, the Windows Firewall, is enabled by default and now provides both inbound and outbound filtering (earlier versions only supported inbound filtering). Outbound filtering can effectively prevent malware from communicating with other computers and fanning out to other systems across the network.

Windows services have always been a favorite malware target: many services are always on and running in a highly-privileged security context (for example, using the local system account (LSA)). In Vista and Windows Server 2008 Microsoft incorporates the notion of restricted services or services that are isolated to a maximum extent. One of the isolation techniques Vista and Windows Server 2008 use is what Microsoft refers to as Session 0 Isolation. Session 0 is the first session that the Windows OS creates when it starts. Session 0 Isolation ensures that only services are allowed to run in session 0. Before Vista and Windows Server 2008 user-level applications could also run in Session 0.

Furthermore, Vista and Windows Server 2008 marks Session 0 as non-interactive, meaning that services cannot directly communicate with users, for example by creating dialog boxes. Vista and Windows Server 2008 services also receive the least possible amount of privileges; what is needed to do their job and nothing less or more: Windows Server 2008 revisited the default permissions and rights that are assigned to services. Vista and Windows Server 2008 services are also constrained in their communications, as Vista assigns a Security Identifier (SID) to each service, implements service-specific access control lists on system resources such as the registry and the file system, and per-service inbound/outbound access restrictions on the Windows Firewall.

Network Access Protection (NAP) is the name of Microsoft's network admission control (NAC) architecture. NAP is a technology that can ensure that only healthy machines connect to an organization's IT infrastructure. "Healthy" in this context refers to: systems that are not infected by malware, that have the latest anti-virus and spyware protection signatures installed, that have the latest security patches installed and that have properly configured security settings. NAP can also require strong user and machine authentication before letting a machine and user onto a corporate network. NAP not only isolates unhealthy and unauthorized machines, it can also heal them, for example, by installing the latest security patches, removing malicious code and/or locking down a system's security settings. The NAP client component is included in Vista and will also be made available for Windows Server 2003 and XP Service Pack 2 (SP2) clients. The NAP server component will be bundled with Windows Server 2008. In September 2006 Microsoft and Cisco jointly announced that they would work on an interoperability architecture for Microsoft NAP and Cisco Network Admission Control (CNAC) - which is an architecture similar to NAP that is built into Cisco network infrastructure products. See the following for more information on this: **www.microsoft.com/technet/community/columns/secmgmt/sm0906.mspx**.

Windows Integrity Controls (WICs) is the name of a new mandatory access control model that Microsoft implements in Windows Vista and Windows Server 2008. Vista and Windows Server 2008 use WICs in addition to the classical Discretionary Access Control (DAC) settings that are based on resource permissions and Access Control Lists (ACLs). WICs are also enforced prior to the classical DAC settings; in other words, WIC settings have precedence over DAC settings. The goal of the WIC model is to block elevation of privilege attacks - these attacks can occur when, for example, a piece of code that is downloaded from the Internet tries to interfere with system resources or processes. Every Windows Vista and Windows Server 2008 system file and process has an Integrity Level (IL) assigned to it in its system ACL. Code and files that are downloaded from the Internet are also automatically assigned an IL. When a process tries to write to a file, Windows Vista and Windows Server 2008 will check whether the process has a higher IL than the file's IL - if it has a lower IL, the process will be blocked from writing to the file. In Vista and Windows Server 2008, browser - downloaded code always gets a low IL, and system files (or files that are owned by the OS) always have a high IL - which makes it impossible for browser-downloaded code to interfere with system files.

Honoring Least Privilege in Windows OSs

User Account Control (UAC) is the least privilege feature bundled with Vista and Windows Server 2008, and is one of the most important architectural security changes in Vista and Windows Server 2008. UAC ensures that any user account that logs on to Windows (even accounts with administrator-level privileges) initially only have plain user privileges. It is only when the user account needs to perform a task that requires administrative privileges that Vista and Windows Server 2008 temporarily expand the account's privileges.

An important UAC property that significantly reduces the Vista and Windows Server 2008 attack surface is User Interface Privilege Isolation. UIPI provides process isolation by ensuring that processes running in the security context of a limited-account user cannot interfere with processes running in the security context of a privileged-account user. UIPI protects against shatter attacks, during which malware that runs in the security context of a limited-account user leverages the Windows inter-process messaging system to inject malicious code into a process that runs in the security context of a privileged-account user.

New and Updated Security Tools in Windows OSs

In Vista and Windows Server 2008 Microsoft enhanced and added a set of important malware protection tools: Windows Defender, the Malicious Software Removal Tool (MSRT), the Security Center, Automatic Updates (AU), and the Internet Explorer (IE) malware protection feature.

Windows Defender is the real-time spyware protection solution that is bundled with Windows Vista. It is the rebranded version of the Giant AntiSpyware solution that Microsoft acquired in 2004. Defender continuously monitors operating system resources such as the registry and the file system that are commonly abused by spyware. If an application attempts to make changes to one of the monitored resources Defender blocks the application and prompts the user to reject or allow the change.

Finally Microsoft added several new malware protection features to Internet Explorer 7 that is bundled with Vista and Windows Server 2008. These features include a phishing filter and better protection against malicious ActiveX controls.

5.4.3.3. Linux

As an open source project, Linux benefits from the contributions of a diverse security community and has a well-deserved reputation for being resistant to attacks and intrusions when configured correctly. Moreover, Linux is unique as a general-purpose operating system that can address multi-level security (MLS) requirements that are traditionally met by military-grade trusted operating systems. HP remains committed to advancing community efforts to enhance Linux security by contributing to the development of MLS features, supporting Common Criteria evaluation efforts, and providing migration services for customers moving their applications from trusted operating systems to enterprise Linux.

The open source development model is frequently credited with strengthening the security features of Linux by virtue of its open review process. Although not all projects enjoy this quality of scrutiny, Linux has received better code and fewer bugs as a consequence. Another result of the active community is a robust set of security mechanisms, cryptographic libraries, and trusted utilities available on Linux for host, network, and application security.

Linux offers a full range of access control mechanisms, including Discretionary Access Control (DAC), Role-based Access Control (RBAC), and Mandatory Access Control (MAC). Supplementing the traditional DAC implementation by the kernel, Linux Security Modules (LSM) is a lightweight framework with hooks in the kernel to enable various access control mechanisms to be loaded as kernel modules.

One such module, initiated by the U.S. National Security Agency (NSA), is Security Enhanced Linux (SELinux), which some Linux distributions now deliver without requiring special setup. SELinux implements a flexible MAC mechanism called type enforcement, which associates each subject and object with a type identifier and allows rules governing type-based access to be defined in a policy file loaded into the kernel at boot time. Because the policy is not hard-coded in the kernel, SELinux provides strong mandatory security in a form that system administrators can adapt to a wide variety of security goals reliably and flexibly. Red Hat Enterprise Linux 5, for example, enables SELinux by default with a MAC policy that provides containment around network-facing daemons. An administrator can deploy a more fine-grained multi-level security scheme by loading a different SELinux policy.

In contrast to the SELinux approach, SUSE Linux Enterprise 10 builds on the inherent security of Linux by integrating a wide range of security capabilities, including encryption, firewalls, certificate creation and management, authentication, access control and proxy management. It is the only Linux distribution to include integrated application-level security with Novell AppArmor. AppArmor tools identify the programs that need containment, capture application behavior in a "learning mode" and turn that behavior into security policy.

HP offers Linux from Red Hat and Novell and supports both Red Hat Enterprise Linux with SELinux and SUSE Linux Enterprise Server with AppArmor. HP has also demonstrated a broader commitment to the development and certification of multi-level security features in Linux. HP has completed three consecutive certifications with Red Hat Enterprise Linux as well as one recent certification with Novell SUSE Linux Enterprise Server.

- HP has completed Labeled Security (LSPP), RBAC, and Controlled Access (CAPP) certification at Evaluation Assurance Level (EAL) 4+ of Red Hat Enterprise Linux 5 on HP Integrity servers, HP ProLiant servers, and selected HP workstations. HP worked closely with Red Hat, NSA, and the security community to enhance multi-level security capabilities, contributing to file system auditing, labeled printing, and labeled networking for greater compatibility with legacy trusted operating systems.
- HP has completed CAPP certification at EAL 3+ of Novell SUSE on HP ProLiant servers, HP Integrity servers, and HP carrier-grade systems. The certification includes systems with the Intel Itanium 2, Intel Xeon, Intel Pentium, and AMD Opteron processor families.

HP continues to enhance the multi-level security capabilities of Linux by supporting the development of auditing, labeled printing, and labeled networking, and submitting these contributions to the community under the Gnu General Public License (GPL). HP also offers a porting kit and migration services to support customers seeking to move their applications from a legacy trusted operating system to enterprise Linux. For more information, see **www.hp.com/go/linuxsecurity**.

5.4.3.4. HP OpenVMS

Security, at its core, is about protecting data and transactions from unauthorized access and ensuring that data is available when businesses need it. HP OpenVMS ships with an out-of-the-box default security architecture that provides "Rings of Protection" (see Figure 4-5, next page) that grant users and applications the least amount of privilege needed to accomplish their tasks. In addition to the rings, which are analogous to multiple layers of physical security found at many buildings, OpenVMS adds yet another dimension of locally exclusive access. A user, process, or device may have access to a particular layer for a specific purpose and still be excluded from access to all other levels for which privileges have not been granted. This provides even stronger assurance against such actions as back-door unauthorized access.

The system itself performs privileged tasks on behalf of a user or application without needing to grant the user that privilege. This design protects OpenVMS from viruses and similar attacks. Data protection extends across the whole system implementation from memory to disk storage to processor to I/O, so that a flexible but secure system can be configured to meet the needs of any enterprise.

Installing OpenVMS without security parameter changes results in a secure environment with no default passwords or accounts with known passwords. One of the first things that OpenVMS requests during installation is the installer's identification and primary security parameters. OpenVMS may be installed without this information, but access is not allowed and installation must begin anew. This default feature is designed specifically to ensure that definitive security precautions are instituted from the very beginning of use.

The resulting system with established rings of security and the ability to monitor and identify users, even those with the most privileges, provides for an implementation of security policy that can be followed directly from the first moment of installation. In addition, OpenVMS provides exceptional data confidentiality (protecting data from unauthorized access) with encryption tools and a default protection scheme that is secure and flexible.

Figure 4-5
OpenVMS rings of protection

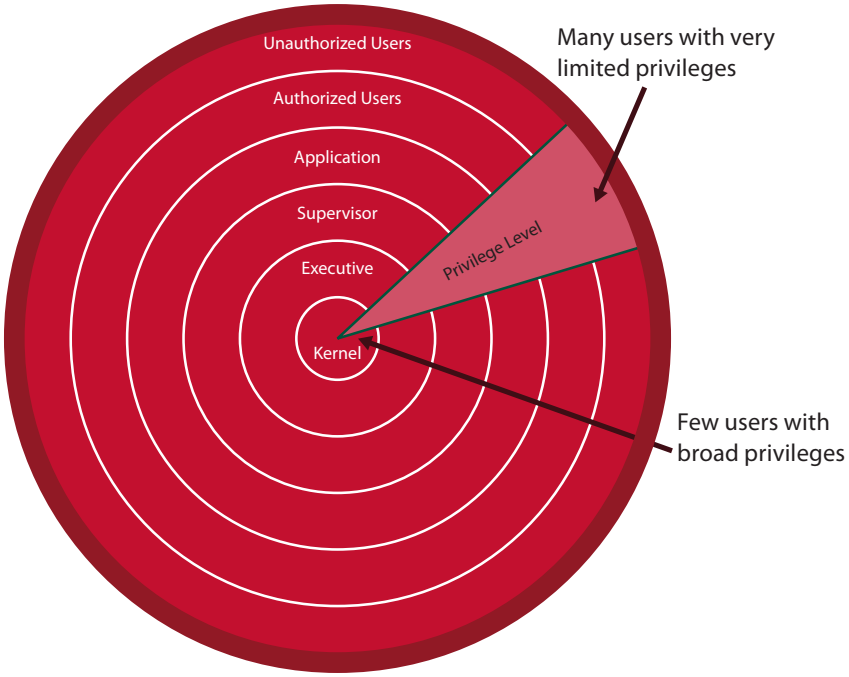

OpenVMS has had security designed in since it was first developed. At the core of the OS is a security model that ensures that every transaction on an OpenVMS system is auditable and access is granted or denied by the security model. The system provides a rich set of tools to control user access to system-controlled data structures and devices that store information. OpenVMS employs a reference monitor concept that mediates all access attempts between subjects (such as user processes) and security-relevant system objects (such as files). OpenVMS also provides a system security audit log file that records the results of all object access attempts. The audit log can also capture information about a wide variety of other security-relevant events.

Because the OpenVMS security model is built into its design, the security features, including the robust logging and auditing functionality, require minimal overhead. As a result, users have the highest level of security in a commercial off-the-shelf operating environment with full performance. When the impact of security is further considered in the context of total cost of ownership (TCO), HP OpenVMS performs very favorably in comparison to competitive environments. The results of a relevant study can be found at **http://h71000.www7.hp.com/openvms/whitepapers/TCS_2004.pdf**.

The OpenVMS security architecture and model apply equally to a single system in a computer lab, or to an entire OpenVMS Disaster Tolerant cluster spread over hundreds of miles. In any case, each access will be tested, audited, and validated. To find out more about OpenVMS security, go to **www.hp.com/go/openvms/security** or review the white paper at **http://h71028.www7.hp.com/ERC/downloads/4AA0-2896ENW.pdf**.

5.4.3.5. HP NonStop Systems

HP NonStop systems provide strong security for a number of financial and other mission-critical applications. With their integrated hardware, software, and middleware, HP Integrity NonStop NS-series and NonStop S-series systems protect your application in these ways:

• Modular operating system: Except for a small kernel, most of the HP NonStop operating system functionality is handled by specialized system processes, such as the memory manager and disk access manager that communicate through interprocess messages.
• Minimum privilege: Not all system or application processes need administrator or root privileges to run, but may be started with the minimum authority required by the customer.
• Processes that run in their own virtual address space: No matter what a non-privileged process does, it cannot view or alter the memory of any other process running on the system unless the processes agree to share portions of their memory. Processes normally communicate by sending messages.

HP Safeguard security management software, included as part of the NonStop operating system for Integrity NonStop NS-series servers and available for NonStop S-series servers, implements a finer-grained subject/object access control model than the one provided by basic system security services.

• Authentication: Safeguard software complements and extends the basic security features by adding advanced support for UNIX-type user names and features such as account expiration, temporary access suspension and restoration, password quality, password history, password change intervals, and automatic user account suspension after excessive logon failures.

• Authorization: Safeguard software can improve server availability by reserving resources for critical production applications, ensuring that only authorized clients can access application servers, and protecting critical data from unauthorized or accidental modification. Authorized users can exercise control over objects such as disk files, tape drives, and other processes. You establish the protection of an object by creating one or more access control lists (ACLs) or protection records for it. An ACL contains subjects or groups of subjects (users) and the access that they are permitted to the object.

• Auditing: Safeguard software audits logon attempts, access to objects, and changes to the security settings for those objects, allowing your security administrators to detect unauthorized system access, detect unauthorized security setting changes, and verify that policies are being followed. Security administrators can specify the objects and the types of access to be audited and how much or how little system activity to record for later review. Safeguard software also logs changes made to an object's ACLs. This record can be reviewed by management and auditors to verify that security administrator activity conforms to established management policies. In addition, Safeguard audits changes to its own configuration.

The HP NonStop Security Review Service provides a comprehensive assessment of the security risks to a business's HP NonStop Server with clear, prioritized recommendations to counter those risks.

There are dozens of NonStop system security enhancements available from HP partners. Customers can take advantage of valuable off-the-shelf features such as single sign-on; support for RSA SecureID tokens; graphical interfaces; enhanced logging and reporting; limiting authorization to specific times, locations, and access devices; and granularity to the individual command level of system utilities. Frequent interaction with these partners allows HP to understand what new APIs should be made available to increase the functionality of NonStop system security.

HP NonStop systems use best-practices technology to provide strong authentication, authorization, and privacy in their overall networking design. This includes support of biometrics, tokens, and PINs for authentication. Least-privilege access, role-based security, and subject/object access control models are used for authorization solutions. The HP NonStop Security Review Service provides a comprehensive assessment of the security risks to a business's HP NonStop Server with clear, prioritized recommendations to counter those risks.

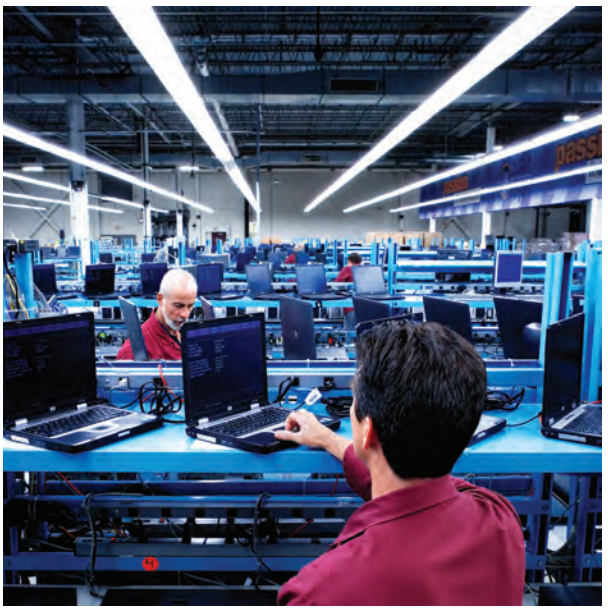

5.4.4. HP Atalla Security Products

HP Atalla Security Products incorporate years of cryptographic expertise and industry best practices into designing and building hardware-based commercial security appliances that meet the high levels of government security requirements.

The Atalla Cryptographic Subsystem (ACS) used in most Atalla products is the first reparable technology to have been validated at Federal Information Processing Standard (FIPS) 140-2 Level 4, the highest government standard for physical security and key management. Where it is appropriate, other Atalla products meet Common Criteria levels 4 and above. Within HP, Atalla is uniquely focused on strong security solutions and cryptographic performance.

HP Atalla products that secure worldwide bank payments networks have set new security, performance, and flexibility standards In the face of increasingly sophisticated threats and escalating risks. HP Atalla is the market and technology leader in strong Automated Teller Machine (ATM), Electronic Funds Transfer (EFT), and Point-of-Sale (POS) network security. Atalla products provide unparalleled performance, price, and protection to over one thousand leading financial institutions, independent software vendors (ISVs), and HP financial industry partners.

The high-performance HP Atalla Ax150 Network Security Processor (NSP), an evolution of the industry-leading Atalla Ax100 NSP family, is a hardened, rack-mountable, 2U appliance with a state-of-the-art tamper-resistant architecture. The Atalla Ax150 NSP provides unrivaled protection for Triple DES and other cryptographic keys when safeguarding value-based transactions. The Atalla Ax150 NSP series consist of three models; each model offers identical security functionality, with different processing capabilities. The high-end HP Atalla A10150 NSP performs up to 950 Triple DES PIN translates per second while using the industry-standard Atalla Key Block. All Atalla Ax150 NSP meet the FIPS 140-2 Level 4 standard.

The HP Atalla Secure Configuration Assistant (SCA) is an easy-to-use handheld device to configure commands, define parameters, and inject cryptographic keys into new-generation HP Atalla network security processors. The SCA is based on a security-enhanced HP iPAQ personal data assistant (PDA), with an easy-to-use graphical user interface that saves time and reduces human error. The Atalla SCA security engine is a custom smart card that performs all cryptographic functions and stores security-relevant data such as cryptographic key components to ensure the highest levels of physical and logical security.

The Atalla SCA is physically secured with tamper-evident seals and has exceptional logical security features. The Atalla SCA application starts automatically on SCA power-on, accepts digitally signed upgrades only from HP Atalla Security Products, and is locked to prevent the installation of any rogue applications. The custom SCA smart card Is certified to the FIPS 140-2 Level 3 standard. Together, the Atalla NSP and Atalla SCA form the only end-to-end truly secure key initialization, configuration, and key management solution on the market.

The **Atalla Key Block** (AKB) is an extensible, secure, industry-standard foundation for cryptographic key management. The Atalla Key Block is the new-generation key management solution from HP Atalla, designed from the ground up to provide unrivalled logical security for Triple DES and other advanced cryptographic keys. No matter what strength cryptographic algorithm is in use, secure key management is pivotal to its effectiveness.

The Atalla Key Block prevents even a knowledgeable attacker from:

- Changing any attribute of any key
- Changing any bits of any key
- Using part of a key as the entire key
- Rearranging any part of a key
- Substituting parts of a key into another key
- Identifying weaker keys

The result of more than three decades of cryptographic expertise, the Atalla Key Block has become the industry standard for cryptographic key management as defined by ANSI X9.24-2004, Retail Financial Services Symmetric Key Management Part 1: Using Symmetric Techniques.

Customers deploying the Atalla Key Block are able to derive maximum value from their cryptographic schemes, overcome key security issues associated with mixed encryption environments, and enjoy extensible protection far into the future. Designed for simplicity of use as well as security, the Atalla Key Block is supported by leading financial institutions, ISVs, and HP industry partners with an interest in the security of financial networks.

The HP Atalla Resource Manager (ARM) is a set of software tools that provide flexible and unique ease-of-use features to HP NonStop server users for managing groups of Atalla NSP to optimize the performance and security of cryptographic operations. ARM enables users of two or more Atalla NSP to simplify their environment by managing the cryptographic access to NSP. IT organizations with a complex Atalla NSP environment or experiencing rapid growth can save time and money using ARM to customize their environment to the needs of their financial institution.

The HP Atalla Trusted Print Center (ATPC) brings robust cryptographic security to the printing of PIN mailers and the re-keying of remote ATMs. The ATPC is a cost-effective in-house solution used to generate PINs (personal identification numbers) and print PIN mailers that advise account holders of their PIN in a secure manner. A PIN mailer is a tamper-evident form where the PIN is never visible to anyone until the intended recipient opens the folded secure form. Until the advent of the ATPC, this essential process within financial institutions has been less than secure.

The protection and changing of unique keys within an ATM or POS device is essential to the security of banking networks. The distribution of key components is a logistical nightmare for financial institutions. Like PIN mailers, the ATPC solution generates, encrypts, transmits, decrypts, and prints key components on secure ATPC forms maintaining best security practices such as the dual-control principle. The two technicians servicing the remote cryptographic device open their tamper-evident forms and inject their respective key components, then communicate the completed tasks to the host application for final test of the newly injected key. The sensitive cryptographic key is never "in the clear" from the NSP until used at the ATM or POS device. The ATPC uses a custom Atalla NSP that meets the FIPS 140-2 Level 4 standard.

The Atalla Remote Key feature is a second technology used to remotely key and re-key ATMs and other cryptographic devices. ATMs are not physically secure while being serviced. For example, there is no dual control or split knowledge of encryption keys. In addition, changing ATM keys are estimated to cost up to $400 per ATM change.

The industry is moving to a faster, more accurate process with no human intervention that will reduce overheads and increase the number of key changes per annum. HP Atalla supports initiatives from ATM vendors such as NCR, Diebold, and others that meet recommendations such those of VISA and MasterCard.

For more information about HP Atalla Security Products, see **www.atalla.com**.

5.4.5. HP Trusted Compliance Solution for Energy

The HP Trusted Compliance Solution for Energy (TCS-e) uniquely helps North American utilities define, implement and maintain their security controls while automating these processes to reduce both the initial and the ongoing costs of NERC CIP 002-009 compliance while ensuring the integrity of NERC evidence.

Utilities and their auditors can easily and immediately see who accessed critical infrastructure and when - for compliance, audits, and forensic investigations - with confidence that this data has not been compromised. To maintain the highest levels of reliability to end customers, automatic real-time searches for security signatures immediately alert utilities to potential security-related or natural crisis events.

Workflow management translates NERC CIP evergreen review requirements into automated reminders to responsible reviewers and approvers, ensuring ongoing compliance. Unlike an off-the-shelf tool that can be used to collect the required documentation, TCS-e was purpose-built around NERC CIP requirements integrating data access, security and workflow. TCS-e provides a comprehensive methodology within the help module, developed by a senior industry consultant to help utilities navigate the complex NERC CIP standards as they plan and execute their security compliance efforts.

The Trusted Compliance Solution for Energy is packaged as a 2U-high, hardened appliance with an embedded cryptographic engine, a Trusted Configuration System for assured deployment and management, and three management services:

• The Trusted Compliance Manager (TCM) securely automates the assembly, review, approval, and audit of NERC CIP 002-009 compliance evidence so that a utility senior manager may sign off with confidence. TCM provides a "Dashboard" that depicts the organization's progress to compliance to NERC CIP cyber security standards. TCM automates the secure management of compliance documents with digital signatures and trusted timestamps. When senior managers are asked to sign off on their team's compliance effort, they can do so with confidence.

• The Trusted Log and Analysis Manager (TLAM) securely collects, compresses, and stores log record data in a trusted, purpose-built, replicated repository that compresses data at a ratio of up to forty to one over a relational database management system (RDBMS). Log records are secured with trusted time stamps and digital signature so there is no question just who did what and when. TLAM is a "Flight Data Recorder" for internal or external audits and forensic investigations. Sophisticated report and query tools scan the still-compressed data repository to return information orders of magnitude faster than a traditional RDBMS.

• The Trusted Real-time Alert Manager (TRAM) is a Motion Sensor that scans log record data from numerous sources, in realtime for potential security-related or natural events against control systems, and alerts trained personnel for further investigation and action. TRAM can detect simple threshold events or complex series of events that may be unique to the control system environment.TCS-e leverages unique security capabilities from HP that provide the trust and reliability found in the global bank payments interchange network. With Federal Information Processing Standard (FIPS) 140-2 Level 4 protection, TCS-e meets the most rigorous U.S. government standards to protect a utility's most sensitive compliance, information technology, and process control data and log records.

5.4.6. Vulnerability Assessment Tools from HP

HP Application Security Center software solutions help security professionals, developers and QA teams save time and money by catching security defects as early in the application development lifecycle as possible.

HP Application Security Center software is designed with flexibility in mind. Some development and QA organizations want to deploy software that is integrated into the development and testing environments. Others want a centralized solution for authorized team members to conduct security tests as needed. Many organizations implement a combined approach in which security professionals manage the overall security program, working with developers, QA teams and security experts. They need flexible solutions to define and manage web application security processes.

HP Application Security Center's DevInspect, QAInspect and WebInspect softwares are designed for developers, QA professionals and security professionals respectively. HP Assessment Management Platform software brings these products together and can be leveraged by each audience for different purposes. When used together, these products provide an effective end-to-end security testing solution for your enterprise.

The HP Application Security Center software provides common security policy definitions, centralized permissions control and web access to security information, and it supports the complete application lifecycle from development to production.

5.4.6.1. Vulnerability Assessment Tools for Security Professionals

Security professionals must secure enterprise web applications and reduce the risk of malicious attacks from hackers. Hackers are constantly finding new ways around traditional defenses in order to break into web applications and web services. While protecting assets and maintaining security awareness in this complex, fast-changing environment, security professionals must also demonstrate the state of your web security and regulatory compliance.

Security professionals must also address an overwhelming number of applications, vulnerabilities and people around the world. They must identify critical applications, maintain a holistic risk management view and give numerous stakeholders visibility into the state of application security across the enterprise. They must scale their assessment processes across the enterprise and throughout the lifecycle to developers, QA teams, other security professionals and business managers who own the applications. Many organizations are striving for proactive application security programs that find vulnerabilities early in the lifecycle to avoid the excessive costs associated with fixing defects in production applications. The security professionals driving these programs need sophisticated software to help them coordinate a global team of people and manage and mitigate application risk. HP offers two software products that can be used separately or together to test web applications and manage your overall security program: HP WebInspect and the HP Assessment Management Platform.

HP WebInspect

HP WebInspect is easy-to-use, extensible and accurate web application security assessment software. Many security professionals begin their application security testing programs with HP WebInspect, which enables both security experts and security novices to identify critical, high-risk security vulnerabilities in web applications and web services. HP WebInspect addresses the complexity of Web 2.0 and identifies vulnerabilities that are undetectable by traditional scanners. HP WebInspect supports today's most complex web application technologies with breakthrough testing innovations, including simultaneous crawl and audit (SCA) and concurrent application scanning, resulting in fast and accurate automated web application security.

HP Assessment Management Platform

HP Assessment Management Platform fully addresses the complexities of today's web application security programs. After using HP WebInspect for a short time, security professionals often need to scale their program to test additional web applications and perform tests more frequently. They need both automated and scheduled penetration testing and more manual expert tests. They need to extend security testing to additional security professionals around the globe as well as developers and QA teams, who address security early in the application lifecycle.

HP Assessment Management Platform supports an advanced global security program that allows multiple participants to get the application security information they need and participate in the assessment and remediation process, while letting security professionals maintain centralized control. HP Assessment Management Platform is distributed and scalable. It provides a web-based interface for a consolidated global view, supporting multi-user lifecycle collaboration and control of application security risk throughout the enterprise. Developers, QA teams and security professionals can use HP Assessment Management Platform as a black box assessment tool across the enterprise to target vulnerabilities that hackers can exploit.

5.4.6.2. Vulnerability Assessment Tools for Developers

Developers are increasingly using products to help them code more securely. Developers know that security defects are like other defects. Catching security defects early eliminates the time and expense associated with later-stage patches. Today's global organizations have thousands of developers dispersed all over the world. In many cases, development is outsourced to third party vendors. Establishing common practices and tools for secure coding is an ever-present challenge. HP offers two products that present different options for developers to test their web applications for security: HP DevInspect and the HP Assessment Management Platform.

HP DevInspect

HP DevInspect simplifies security for developers by automatically finding and fixing application security defects. HP DevInspect also helps developers build secure web applications and web services quickly and easily, without affecting schedules or requiring security expertise. HP DevInspect is installed on an individual developer's system and, using a Hybrid Analysis approach, combines source code analysis with black box testing to reduce false positives and find additional security defects.

HP DevInspect integrates with the following integrated development environments (IDEs):

- Microsoft Visual Studio 2003 and 2005
- IBM Rational Application Developer 6 and 7
- Eclipse 3.1 or higher

HP DevInspect supports C#, Java, Visual Basic, HTML, XML, SOAP, WSDL, JavaScript and VBScript.

HP Assessment Management Platform

Many development organizations also use HP Assessment Management Platform, which developers can use to conduct assessments of their code as needed. Developers can use HP Assessment Management Platform to conduct black box testing of their application, targeting only exploitable security defects. HP Assessment Management Platform conducts comprehensive tests for all web applications, regardless of the language in which they are built. HP Assessment Management Platform also includes flexible reporting capabilities that let development teams share information and security policies with QA teams and security professionals.

5.4.6.3. Vulnerability Assessment Tools for QA Professionals

QA teams use security products to help them find security defects in web applications. Security testers have always focused on functionality and performance. Now that web applications are maturing, QA teams are conducting more focused and comprehensive security testing on web applications.

HP QAInspect

HP QAInspect applies innovative techniques to identify security defects from the hacker's perspective. HP QAInspect reports on vulnerabilities with detailed security knowledge in a way that QA professionals can understand with a concise, prioritized list of vulnerabilities and thorough vulnerability descriptions. Analysis results yield detailed information on possible types of attacks, including cross-site scripting (XSS) or SQL injection, as well as on compliance issues related to regulations, such as SOX, HIPAA and PCI.

HP QAInspect integrates with the following testing tools:

- HP Quality Center software
- HP WinRunner software and HP QuickTest Professional software
- HP Business Process Testing software
- IBM Rational Software Delivery Platform (SDP), Rational ClearQuest and Rational Functional Tester

HP Assessment Management Platform

Many QA teams also use HP Assessment Management Platform to assess their applications. HP Assessment Management Platform conducts comprehensive tests for all web applications, and its automatic scheduling capability lets QA teams schedule regular web security tests. HP Assessment Management Platform also includes comprehensive reporting capabilities that help QA teams share information and security policies with development teams and security professionals.

5.4.7. HP Enterprise Mobility Suite

The HP Enterprise Mobility Suite (HP EMS) provides enterprises with a secure and cost- effective solution to deploy mobile devices to the field while ensuring manageability and security.

HP EMS enhances the security of sensitive corporate data by extending corporate security policy enforcement to the mobile device and by reducing security vulnerabilities with Over-The-Air (OTA) access to the entire fleet of devices. HP EMS also provides a dedicated enterprise management solution for automating mobile device management, including OTA device setup, diagnostics and application management.

HP EMS gives enterprises the following security and policy conformance mechanisms for their mobile devices:

- Remote lock or wipe of compromised devices. Enterprises can remotely push lock or wipe commands to lost or stolen devices to restrict access to data and device functionality. Locking a device immediately protects data until the device has been recovered or wiped.
- Device lockdown. EMS supports the native Windows Mobile lockdown features with policies to disable Bluetooth, Wi-Fi, IR, camera, and removable storage.
- Policy conformance. Enterprises can set IT policies for devices to regularly self-audit these devices for conformance to corporate policies, such as those requiring encryption and power-on passwords. They can flag non-conforming devices and automatically apply policies so devices return to conformance immediately.
- Support for user certificates (for HP iPAQ only). User certificates can enable secure mobile access to the network behind the firewall.

HP EMS can manage both mobile phones and non-phone devices (for example: PDAs) over any available HTTPs connection including cellular, WIFI, or tethered connections. It leverages the mobile- optimized Open Mobile Alliance (OMA) - Device Management (OMA-DM) standard for supporting multiple device platforms.

For more information on the HP Enterprise Mobility Suite see **www.hp.com/go/ems**.

5.5. Host Security Summary

Host security has traditionally been a military-grade solution with high costs in the areas of user satisfaction, user productivity, and operations - in addition to the cost of the solutions themselves. Host security is transforming to meet the needs of businesses and other organizations, which are driving secure hosts to deliver ease of administration, flexible role-based access control, useful privilege management, and security balanced with performance.

The concept of relying solely on a bulletproof perimeter defense is evolving into the concept of layered defenses that acknowledge the real threat environment. Furthermore, the layers need to extend all the way down to the servers themselves. The motivation for these changes comes from far-reaching, global virus attacks, such as the Blaster and Sasser worms, that have easily crossed secured perimeters.

HP has intently examined the issues related to host security to enhance the delivery of platform security through the operating environment. The results are new tools and techniques that reduce the risk to enterprises without ballooning TCO or creating an unacceptable customer or user experience.

6. Network Security

The enterprise network connects all other trusted infrastructure elements. A properly secured network protects and integrates its hosts, while remaining functional in the face of business-driven change and today's countless threats to information availability, integrity, and confidentiality. This section focuses on data network security and limits its coverage to IP-based networks. It discusses network security threats, defenses and design, and the selection of specific network security components.

6.1. Environment

Enterprise networks are changing. Modern networks have a diversity of components with varying trust levels; they are no longer simply fortresses encircled by defensive rings. Traditional enterprise networks have an internal compartment devoted to internal communications and a carefully isolated compartment - commonly called a demilitarized zone (DMZ) - devoted exclusively to externally accessible services. Firewalls control access between the internal network and the DMZ, and between the entire network and the Internet.Three factors have caused the enterprise network to change dramatically: Internet delivery of services, telecommuters and mobile workers, and outsourcing. These factors are illustrated in Figure 4-6.

Figure 4-6
Enterprise network trends

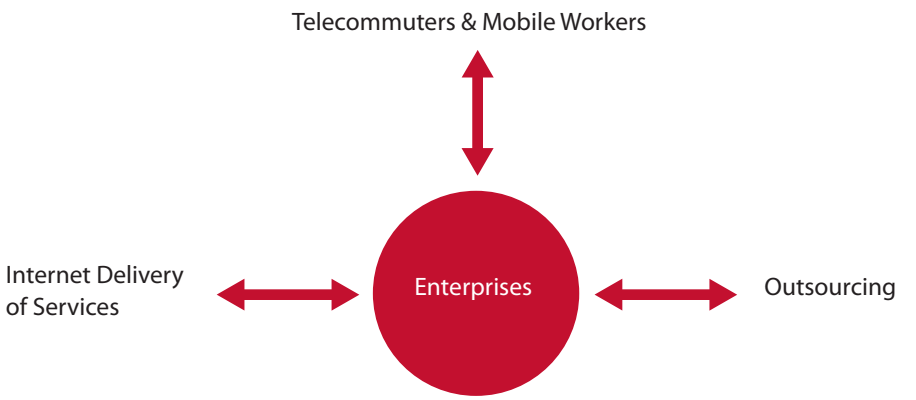

6.1.1. Internet Delivery of Services

Organizations use the Internet to deliver increasingly complex services in intra- and inter-domains, including collaboration and transactions with vendors, customers, and partners. This process often relies on interactions between groups of systems on different enterprise networks. Therefore, external services hosted in the DMZ interact with internal systems in increasingly complex ways, complicating the relationship between the DMZ and the internal network.

6.1.2. Telecommuters and Mobile Workers

Organizations rely on telecommuters and mobile workers to perform critical tasks that require access to internal applications and data. Organizations must provide access to internal resources from almost any location and for a variety of devices over which they have varying degrees of control. These requirements further blur the network perimeter.

6.1.3. Outsourcing

Organizations are distinguishing between their core competencies and other business-critical activities in order to better compete in the global economy. Many organizations are aggressively outsourcing critical work to distant partners or delivering global services based on their core competencies. IT has responded by making internal applications available externally via virtual private networks (VPNs), leased connections, terminal serving, and reverse proxies.

Taken together, these three trends result in networks with a variety of users, segments, and hosts that are authorized to do different things and are trusted at different levels. Fortunately, while networks have become more complex, their security capabilities have become more sophisticated and autonomous.

6.2. Network Security Analysis and Planning

Network security addresses enterprise network technologies that connect to the Internet and to extranets in addition to the boundary of the IT infrastructure. Network security must also focus on the security of wireless networks. Controlling access to network resources and providing lower-level prevention and detection of attacks allows enterprises to optimally protect their information assets.

6.2.1. Approach

A traditional approach to network security implementations is to encircle an unsecured network with a perimeter defense solution that controls access to the network. Perimeter defense is an integral part of an overall defense strategy. However, within the perimeter, a user left unrestricted may cause intentional or accidental damage. The network can be extremely vulnerable to a hostile party gaining access to a system or application inside the perimeter, and it can be compromised by an authorized user. Crucial steps include ensuring that all devices on the network are authorized to connect, the devices are up to date in configuration according to organizational policy, and the network is adequately defended against attack.

Deploying a collection of security techniques and tools, including firewalls, intrusion detection systems (IDSs), intrusion prevention systems (IPSs), and VPNs, can help enterprises to ensure overall network security. However, as network boundaries expand and fluctuate due to interconnected business relationships with employees, vendors, and customers, efficiently making and enforcing access decisions becomes challenging. While protecting the network is a high priority, minimizing friction to legitimate authorized usage is the goal. Even the most trusted and secure systems are exposed to a variety of threats when employees take them offsite. The health of the devices must be reevaluated every time they attempt to connect to the network. As in all other areas of network security, nothing is static and nothing should be assumed or taken for granted.

There are a number of approaches to network security, and there is not a single solution that applies to all organizations. Many factors impact assessing and developing a strategy for authentication and authorization of devices and users, and for assessing and remediating threats posed by devices attempting to access the network. Ultimately, any network security solution, no matter how sophisticated or capable, does not mitigate or eliminate the ultimate responsibility of the organization. Protecting sensitive data where it is stored and exerting appropriate controls for when and under what terms data is accessed are critical components of an organization's complete security solution.

From a technology perspective, network security is a rapidly developing area, with new threat vectors and new security solutions being developed and refined constantly. Organizations rarely have the necessary internal resources to best evaluate, determine, and implement appropriate courses of action. Given the breadth and heterogeneity of the HP portfolio, HP consultants are uniquely qualified to advise organizations about how best to design and implement a network security architecture that supports business goals.

HP Services employs a proven solution framework to ensure that networking technology is aligned with business strategy to become a business enabler, not an obstacle. The process begins with an analysis of the organization's business requirements, including business drivers and associated metrics. It includes economic analysis of the impact of technology on an organization's operations and compliance with governance and business objectives.

6.2.2. Understanding Security Risks and Threats

Evolving regulatory and legal requirements are increasing enterprise risk exposure to a level where IT risk management should be a top priority. Network administrators face a witch's brew of dangers: vulnerability scanning; DoS attacks; hijacking of networks to do harm elsewhere; defacement of public web sites; physical intrusion into sensitive areas; abuse of kiosks, hotspots, and other public computing facilities; wide distribution of high-quality attack tools; network mapping and port scanning; vulnerability scanning; war dialing; and war driving. The list seems endless.

As stated previously, the dangers do not always come from the outside. External threats can also be realized by internal attackers who may be employees or contractors engaged by the enterprise. Another area of vulnerability is a trusted network connection, such as a connection with a vendor or trading partner that has experienced a network security breach. Due to the wide availability of ready-to-run attack software tools, attacks can even be mounted by unsophisticated users - sometimes referred to as script kiddies.

There are four general categories of security threats to the network:

1. Unstructured threats consist of random attackers using a variety of tools to attempt to crack protected systems. The tools used include password crackers, credit card number generators, and malicious shell scripts, among others.

2. Structured threats are usually generated by technically competent individuals or organizations. They seek to obtain access to highly sensitive data, and their attacks include development of sophisticated attack plans. They are often sponsored by organized crime and well-financed organizations.

3. External threats include structured and unstructured attacks. They may be random errors or attacks with malicious or destructive intent.

4. Internal threats usually involve disgruntled or former employees. These threats seem the most ominous, but measures are available to mitigate them. Internal threats may result from user ignorance, a knowing violation of security policies, access of malicious web sites, or a download or received e-mail that contains viruses, worms, spyware, or other malware.

6.2.3. Understanding Types of Attacks

The security threat environment is growing more unpredictable, and new threats are emerging and evolving at an increasingly rapid pace.

The best strategy for protecting the enterprise is to adopt a security posture that is as proactive as possible. This section discusses types of security attack categories and specific mitigation methods.

Attacks that compromise resources consist of four basic categories:

1. Reconnaissance attacks occur when an attacker attempts to discover and map systems, services, and vulnerabilities. Typical tools and techniques include packet sniffers, port scans, ping sweeps, Internet information queries, and vulnerability scanning software.

2. Access attacks occur when an attacker attempts to retrieve data, gain access, or escalate access privileges.

3. DoS attacks occur when an attacker attempts to disrupt the service that a resource normally provides.

4. Worm, virus, and Trojan horse attacks occur when an attacker attempts to damage or corrupt a system, replicate malicious code, or deny services or access to network resources.

6.2.3.1. Reconnaissance Attacks

Reconnaissance attacks are performed with packet sniffers, port scans, ping sweeps, Internet information queries, or vulnerability scanning software.

Packet Sniffers

A packet sniffer captures data that is transmitted in cleartext on the network. Examples include user names and passwords transmitted in applications such as telnet, FTP, and e-mail. Detecting the sniffer is difficult unless direct access is available to the system running the sniffer.

Mitigation strategies for packet sniffers include:

- Secure password mechanisms thwart packet sniffers from capturing user names and passwords. The options in this area include one-time passwords or encrypting the authentication handshake between a client and a server. Because they are typically good for a short time period, such as a minute, even one-time passwords should not be transmitted in the clear whenever possible.
- Anti-sniffer tools detect the presence of a sniffer on the network. They must be in place for a period of time in order to detect anomalies that occur when an unauthorized sniffer is launched on the network.
- Switched network infrastructures greatly reduce the effectiveness of packet sniffers in the enterprise.
- Cryptographically secure channels for transmitting data are the best way to render a packet sniffer irrelevant.

Port Scans and Ping Sweeps

Port scans and ping sweeps cannot be prevented entirely. IDS systems at the network boundary and on the host can detect these types of attacks and notify the administrator that an attack is underway.

Internet Information Queries

Domain Name System (DNS) queries can reveal the IP addresses of systems on a network. This can be very useful for IT personnel to manage the network. On the other hand, an attacker can use the IP addresses to launch a ping sweep to map the network, and then a port scanner can be used to provide a list of all services running on the network.

Vulnerability Scanning Software

These tools are typically intended to enable IT personnel to efficiently find vulnerabilities such as permissively configured hosts, missing patches, and weak passwords. In the hands of an attacker, however, they can point to a successful attack.

Intrusion Detection Systems (IDSs)

IDSs can detect patterns of activity associated with vulnerability scanning. In addition, internal use can be controlled by clear and well-enforced policies: only security personnel should be authorized to use vulnerability-scanning software. Of course, properly patched and configured systems also play a key role.

6.2.3.2. Access Attacks

Access attacks can take the form of password attacks and trust exploitation attacks. Network access control technologies are particularly important tools for defending against access attacks. They are a key element of an overall network security architecture customized for an organization's business environments and goals. Network access control technologies are discussed in more detail later in this chapter.

Password Attacks

Password attacks are executed by malicious users in order to retrieve data or escalate privileges. They are mitigated as follows:

- Use strong passwords. Characteristics of strong passwords include at least eight characters, upper- and lower-case characters, numbers, and special characters. Password management software can require strong passwords. A key element is training users and enforcing password policies, for example, forbidding employees to keep passwords on sticky notes at their desks.
- Expire passwords regularly. Password expiration periods depend on the business risks associated with unauthorized access to the protected data or systems, the likelihood of password compromise, and the expected frequency of password use.
- Disable accounts. After a specific number of unsuccessful password attempts, the user account should be disabled.
- Do not transmit plaintext, static passwords. Use one-time passwords or encrypted authentication credentials.

Trust Exploitation Attacks

Trust exploitation attacks involve a user or system taking advantage of privileges that a system has granted (either to all users or to specific users) without an appropriate level of authentication. Systems on the outside of a firewall should never be entirely trusted on the inside of the network. All too often, systems or network administrators establish trust between a user and some data based solely on an IP address. For example, a network administrator might allow access to an internal website from the Internet based on an IP address at a user's house. This is an insecure access method because an attacker could use (spoof) the same IP address.

Examples of trust exploitation attacks include man-in-the-middle and port redirection. In a man-in-the-middle attack, the attacker becomes an intermediary in a communication session between two nodes in order to capture or alter information. Port redirection works by compromising a target system to listen on a certain configured port and redirect all packets to a secondary destination.

Trust exploitation is mitigated by preventing trust between external hosts and internal hosts, properly authenticating users, and using secure protocols for sensitive communication sessions.

6.2.3.3. Denial-of-Service (DoS) Attacks

DoS attacks are defined simply by their name: the attacker denies a particular service that is normally available to users. It is important to note the method that an attacker uses to execute a DoS attack. The most common type of DoS attack is a Distributed DoS (DDoS). This type of attack is executed through the distribution of malicious code to a large number of systems. The most common delivery methods are distributing e-mail attachments and exploiting target systems in order to deposit the DDoS code.

DoS and DDoS attack mitigation is straightforward, but it is difficult to completely eliminate vulnerability to DDoS attacks. Mitigation includes proper anti-spoofing configuration of routers and firewalls and the use of anti-DoS features on routers, firewalls, and hosts.

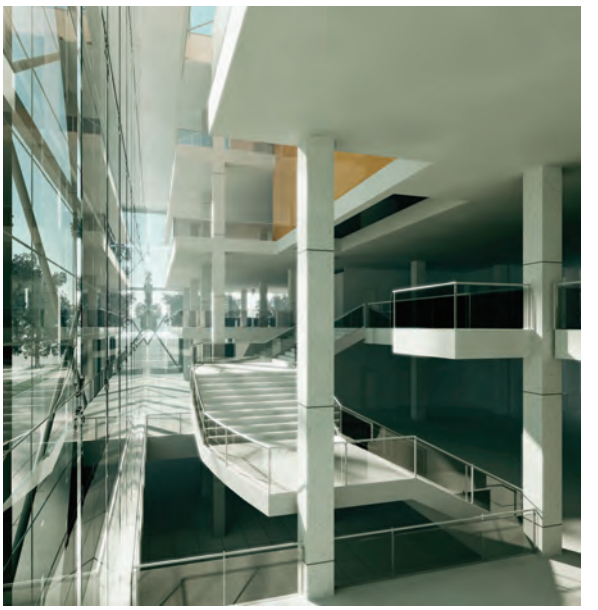

6.2.3.4. Worm, Virus, and Trojan Horse Attacks

A worm executes arbitrary and often malicious code on a host, copies itself to the system's memory, and then copies itself to other computers. A virus is a piece of software attached to another program. When the user's normal program launches, the virus executes and causes unwanted or malicious actions on the host computer. A Trojan horse is different in that it is written to appear entirely benign. However, it executes malicious activities on the host computer.

Mitigation of worm, virus, and Trojan horse attacks is fairly straightforward, but in a large enterprise it can be a challenging task. Mitigation is accomplished through properly using anti-virus software and updating anti-virus signatures. Effective use of anti-virus software also includes installing it on enterprise servers - especially mail servers because e-mail is a significant delivery method for these attacks. As with access attacks discussed previously, network access control technologies are important tools in defending against threats caused by the spread of worms and viruses. They provide mechanisms to enforce the compliance of endpoints for admission to the network environment.

6.3. Principles of Design

Network security is an exacting discipline, and successful implementation requires significant attention to detail. Moreover, with so many variables and available options, it is essential to keep in mind one key point: an effective security solution reflects and protects key business processes and goals. Business drivers must remain foremost, with the network security architecture design supporting them. Identifying the business problems to address through the security solution is a critical element of the design process.

The tools and technologies employed are simply measures to solve business problems. This important step is often performed with the assistance of a knowledgeable and skilled neutral third party able to objectively and effectively assess all of the relevant factors in context, and to put them in the appropriate perspective. Measures that harden and secure the network against internal and external attacks include:

- Compartmentalizing the network (logically, not physically or geographically) to group resources that require similar levels of protection and promote efficient management. This applies not only to legacy systems but also to the overall network architecture. The business drivers should be the key element in determining compartmentalization methods.
- Deploying firewalls at the network level (via network hardware devices, servers, or other products) and on individual workstations and devices.
- Hardening the TCP/IP stack with restrictive settings.
- Deploying port and packet filtering features built into OSs.
- Conducting ongoing training and user awareness - including internal users, vendors, and partners - to ensure continued compliance with organizational security policies.

As with any complex endeavor, however, the right overall direction is essential. A few vitally important principles stand out for any network type and use. These principles should orient the design effort - along with good security policies, mitigation strategies, techniques, and tools. Table 4-2 presents an overview of key design principles. More detail about each design principle can be found in Appendix A.

Table 4-2
Key principles of design for network security

Type	Key Points	Benefits	Cautions
Standardization	• Reduces complexity • Deploys widely tested and trusted tools throughout the enterprise • Benefits from balancing with diverse architectures • Conserves resources	• Reduces the number of threats, risks, and vulnerabilities to identify • Reduces the number of countermeasures to implement • Uses widely tested and trusted approaches	• Does not prevent a single successful attack from being replicated widely • Requires some diversity for a layered, resilient defense
Least Privilege Access	• Requires robust means of establishing and managing digital identities • Grants minimum access to systems or networks needed for business requirements	• Prevents unnecessary resource access • Mitigates the associated risks of resource misuse	• Requires sophisticated tools and processes for user privilege administration • Requires a greater financial ? investment
Layered Defense	• Spans physical, technical, and administrative security measures • Limits risk by combining countermeasures	• Protects the enterprise with multiple forms of defense	• Makes networks more complex and expensive
Redundancy	• Requires justification by business needs	• Enables networks to withstand failure of individual components • Enables networks to withstand successful attacks on individual components	• Makes networks more complex and expensive
Compartmentalization	• Uses logical, not physical, compartments • Works with geographically dispersed systems • Facilitates layered defense and standardization; flexible and adaptable	• Defines access policies centrally and implements them at compartmental boundaries • Accepts changes in business structure and operations without requiring changes to the physical network • Contains security breaches to mitigate damage to the overall network	• Requires careful design • Requires increased cost and complexity with increasing number of compartments

6.4. Securing Network Perimeters and Managing Network Access

Enterprises typically have a multi-level security structure. The first level is the perimeter of the corporate network. To reduce the threat of industrial espionage or deliberate sabotage, only employees, authorized contractors, or other business partners (via an extranet) are allowed access. Although this safety net is difficult to enforce entirely, it thwarts the attempts of casual attackers and creates an obstacle for sophisticated intruders. What does this mean for wireless implementations? First, the secured perimeter must be accessible to mobile devices. Second, information used to access the perimeter from mobile devices must be encrypted to ensure that it is not intercepted or falsified. Typically, the solution to both of these problems is a VPN.

The most sensitive applications need to maintain an additional level of security configuration that includes authentication, authorization, and auditing (AAA). Users with a business need for accessing the application must authenticate to the application first. Depending on their roles and responsibilities, users can have different authorization levels (read-only, modify, delete) or authorization to different subsets of data. A log preserves an audit trail of all actions requested and performed.

Organizations should consider these and other business needs when determining how to secure the network perimeter and manage network access. Key steps in this process include defining security requirements, designing the perimeter, determining access control techniques, and conducting ongoing evaluation and assessment.

6.4.1. Defining Security Requirements

Prioritizing assets in terms of their importance and vulnerability helps organizations identify security needs based on the requirements of the resources being protected.

The result is an informed decision as to the appropriate level of expenditure (both human and financial) for resource protection.

Outward-facing resources such as public web sites need to be accessible but also protected against intrusion. At the other end of the spectrum are resources such as financial systems and sensitive proprietary data that require the most restrictive controls. Occupying the middle ground are resources such as extranets or other sites made available to certain authorized external users (such as vendors, customers, or partners) under specific, well-defined circumstances.

An effective design presumes a solid understanding of the business model and goals and how they are instantiated throughout the enterprise in various systems and resources. Also required is a clear picture of which users and systems - internal and external - require specific access to each resource or compartment. Based on this analysis, which requires continual updating as the business and operating environments evolve, a picture of the threat environment emerges and drives the design of effective and appropriate perimeter controls.

6.4.2. Designing the Perimeter

Perimeter design should be based on a clear understanding of security requirements and the likely threats to resources. It must also include the ability to adapt as protected resources and the threat environment change. The concept of the network edge is subject to many interpretations, particularly in a complex business environment that constantly shifts in response to changes in external relationships (with vendors, customers, or partners). Regulatory compliance constraints can also be a factor as businesses move to a global platform and seek compliance with a variety of foreign regulatory environments.

More specifically, what type of access controls are called for at the perimeter? An example of a common challenge is an unauthorized wireless AP deployed by individual users or workgroups. With the advent of inexpensive and simple wireless gear available at the consumer level, any Ethernet jack becomes a potential unsecured entry point into the network. The devices' default configuration is usually set for unrestricted access, and most users do not know how to properly secure them. This is a prime example of why ongoing user training and awareness are critical. When users perceive a need for wireless access, they should understand enterprise policies and work with the IT department.

Compartmentalization is an architectural approach that is gaining wide acceptance for a network's security architecture implementation. It enables grouping of resources requiring similar levels of protection in such a way as to provide effective and efficient security management and monitoring. This is a key element in HP's Adaptive Network Architecture (ANA), which is described in more detail later in this chapter. Network compartmentalization lets organizations define multiple logical perimeters - each with its own security policies and security architecture - reflecting the specific requirements of the resources contained within each compartment.

Still, the first security measure in most enterprises is a perimeter defense. Enterprises have become very dependent on firewalls and other perimeter protection systems to safeguard their networks. Because it is difficult to secure all the systems in an enterprise and keep them secure, it is necessary to rely on the perimeter as the first line of defense.

Traditionally, firewalls have been mistakenly viewed as magic black boxes; enterprises tend to install firewalls and forget about them. Perimeter security planning and design should begin with risk assessment and consider perimeter defense strategies and standards such as defense-in-depth (layered defenses), trust zones, and hardened systems. Perimeter security implementation should encompass routers, dial-up modems, switches, wireless networking, firewalls, IDSs/IPSs, VPNs, and ongoing network security assessments.

6.4.3. Determining Access Control Techniques

With user devices such as notebooks, PDAs, and smartphones frequently leaving the known safe environment of the enterprise network for remote use, nothing can be assumed about the state of the device. Any device attempting to access the network must be confirmed safe or quarantined for remediation. Techniques for remediation often include evaluation and updates for OSs, applications, anti-virus software, spyware software, and firewalls.

Access control decisions in a modern enterprise network must be more granular than a simplistic yes or no, permitted or denied. Assessing and remediating endpoint compliance is a key part of the system. Sound network design includes AAA capabilities, and none of the approaches for protecting and defending the network perimeter reduces the clear requirement to protect data where it resides. Even if unauthorized or ill-intentioned users gain network access, protected data is safe from unauthorized access or tampering.

A device that is malicious by design or incorrectly configured can pose a substantial threat to an enterprise network. A good proactive security posture dictates that whenever a device attempts to connect to the network, there must be a decision about whether, and under what terms, to allow or deny access. In turn, this access decision must be enforced.

A multi-faceted system is an integral part of the overall network design. A variety of access technologies is available, and selecting the right one (or combination) for a particular environment and situation is essential. VPNs and wireless networking are two key technologies commonly used to manage access. For more information about VPNs and wireless access control technologies, see the Network Security section.

For policy enforcement, the organization needs to apply appropriate access control policies - often integrating authentication, authorization, and access control enforcement. Supporting policy enforcement is a policy decision engine, which evaluates the known attributes of a device, queries for more data as needed, and ultimately comes to an informed access decision.

6.4.4. Conducting Ongoing Evaluation and Assessment

A process should be developed and implemented to measure the effectiveness of security solutions on a regular basis. This typically includes associated metrics for understanding the benefits to the organization derived from implementing the solutions and trends, as well as a process for developing strategies for improvement.

6.5. Securing Wireless Access

In many ways, wireless security is just like wired security and the issues are largely the same. Regardless of the medium, every system needs to safeguard proper authentication, privacy of transmission, prevention of viruses, and protection against DoS attacks. The differences arise from the fact that mobile devices and their transmissions over an unshielded medium (air) are inherently more vulnerable to impersonation, sabotage, and interception.

6.5.1. Securing Wireless Personal Area Networks (WPANs)

Bluetooth is a Wireless Personal Area Network (WPAN) technology intended primarily for cable replacement. It offers simple connectivity between a number of personal devices and corresponding peripherals, such as headsets, phones, printers, keyboards, and pointing devices. It also supports consumer electronics such as cameras and VCRs. Weak protocols can represent a major area of security exposure. Poor stack implementations pose another. However, in the case of Bluetooth, most vulnerabilities stem at least partially from suboptimal device configuration. For example, the default settings of many phones can leave them susceptible to Bluejacking (unwanted messages disguised as business cards). And careless usage may even result in intercepted and manipulated data traffic.

One of the strengths of Bluetooth is that there are several different levels at which the user can specify connectivity options and balance the requirements of usability and security. However, this complexity can also pose a potential risk if users are not familiar with the implications of the settings.

The Bluetooth specifications contain many different security elements and their implementation can entail even more considerations. One logical way to compartmentalize the model is to look at it chronologically, by analyzing each step required to use and operate Bluetooth.

Hardware Selection: The first question to ask is whether a Bluetooth-enabled devices is even needed and if so which device should be selected. The type of device will determine the set of services that might be relevant to it, both as a consumer and provider. But it will also impact the human-machine interface for configuring Bluetooth. So for example, a peripheral such as a Bluetooth headset or GPS receiver has very limited options to set a PIN or configure refined levels of access.

Radio Activation: The second question to ask is whether Bluetooth is always needed. By turning the radio off when it isn't being used it is possible to reduce the attack surface and conserve battery life.

Visibility: The fact that Bluetooth is active does not mean that other users and devices can see it. It may be necessary to make the device visible at times but there is no need for general visibility.

Pairing: The pairing process precedes most active Bluetooth usage. It is necessary so that only devices that belong together can join together. Thus, it is important to be very careful while executing this procedure.

Authorization: After pairing the user may still want an active approval of each connection request. This can be tedious but does provide an additional level of security.

Service Configuration: Even paired devices may not need access to all the services available on a particular device. The user may be able to decide which are required and which are not, as well as what security to enforce on each of the services.

In order to ensure a secure deployment of Bluetooth it is important to evaluate all levels of operation systematically.

6.5.2. Securing Wireless Local Area Networks (WLANs)

There are three acceptable ways to secure Wireless Local Area Networks (WLANs) today: WiFi Protected Access (WPA), VPN with IP security (IPsec), and network access control. Another alternative is to leave wireless access points entirely open and unsecured, but place them outside the perimeter (firewall) of the network. This enables use by vendors and visitors who need wireless access, for example.

6.5.2.1. WiFi Protected Access (WPA)

WPA provides acceptable security using either a Pre-Shared Key (PSK) or in conjunction with the 802.1x protocol and Remote Authentication Dial-in User Service (RADIUS). For enterprises, the 802.1x protocol is advisable because PSK management does not scale easily or provide user-based authentication.

The advantages of this approach are that it involves the least infrastructure and is most efficient in bandwidth. The disadvantage is that all the Network Interface Cards (NICs) and access points (APs) must support WPA. Large implementations may therefore experience significant costs for upgrading and replacing equipment.

6.5.2.2. Virtual Private Network (VPN)

VPNs enable organizations to use the public Internet for secure communications. VPNs provide authentication, confidentiality, and message integrity services that enable organizations to trust information sent over the Internet. The VPN Consortium (VPNC; **www.vpnc.org**), an international trade organization for manufacturers in the VPN market, specifies three VPN technologies:

- IPsec with encryption in either tunnel or transport mode
- IPsec over Layer 2 Tunneling Protocol (L2TP)
- Secure Sockets Layer (SSL) 3.0 or Transport Layer Security (TLS) with encryption

Wireless connectivity can be protected with a VPN. With this technique, all WiFi APs must be placed in a virtual local area network (VLAN) that terminates at a VPN concentrator. Users must authenticate to the VPN to access the network, which also ensures that the content is encrypted. The benefit of a VPN is that it is a proven technology that works with all NICs and APs. The disadvantage is that it adds about 30% of overhead onto the data being transmitted, which constrains air traffic in environments limited by capacity. In addition, the network must be configured to use VLANs and a VPN concentrator must be available.

6.5.2.3. Network Layer Access Control

The 802.1x protocol provides a limited form of network access control to gain access to a network, the user must successfully authenticate. However, no other controls limit user authorization. To address the need for refined access control, new appliance products are appearing in the market. They are known as role-based access controls (RBACs), network access controls, or simply wireless switches.

These appliances require the user to start a browser session before accessing the network. A switch redirects the browser to an authentication page and typically authenticates to a RADIUS server. In addition to simply allowing or rejecting access, they can constrain access based on any combination of time of day, user group, location (VLAN), target subnet, and target application. For example, employees in the finance department may be allowed to access the payroll system from their offices at any time. From the lobby, access to the payroll system could be restricted to office hours and access to e-mail could be allowed at any time.

Beyond the virtually unlimited combination of access control rules, these products can offer enhanced services such as bandwidth throttling and over-the-air encryption. They are similar to a VPN. In some respects, they are a VPN superset because a VPN client operates between the client and the switch. Network access control is NIC/AP independent but requires dedicated equipment, which can be expensive. It is ideal in environments with highly diverse user groups for limiting access to the resources and applications that each group needs. Some of the main customer segments include educational institutions and airports. Universities have diverse needs for groups such as professors, administrators, students, visiting professors, and guests. Airports have diverse needs for venue operators, passengers, airlines, airport security, and baggage handling.

Table 4-3
Wireless LAN security methods

Technology	Advantages	Disadvantages
WiFi Protected Access (WAP)	• Little additional infrastructure • Efficient bandwidth	• NICs and APs need to support WPA (adds cost to upgrade and replace equipment)
VPN (IPsec)	• Proven technology • Works with all NICs and APs	• Adds 30% overhead • Can constrain air traffic
Network access control	• NIC/AP independent • Ideal for highly diverse sets of users	• Requires dedicated equipment (adds cost)

6.5.2.4. Wireless Wide Area Networks (WWANs)

Today's Wireless Wide Area Networks (WWANs) include security provisions that are enforced by mobile terminals and base stations. Current examples of WWANs include Global System for Mobile Communications (GSM), General Packet Radio Service (GPRS), High-speed Download Packet Access (HSDPA), IS-95 Code Division Multiple Access (CDMA), IS-136 Time Division Multiple Access (TDMA), and Single Carrier Radio Transmission Technology (1xRTT).

Although encryption algorithms provide an effective barrier to the vast majority of attackers, it is important to realize that they, like virtually all other encryption methods, are not uncrackable. Both the CDMA and GSM algorithms are reported to have been cracked. The value of the protection does not lie in providing a completely secure environment for sensitive transactions. Instead, it offers an obstacle so that monitoring and interception of random or bulk transmissions is simply not cost-effective or easy for a casual or unsophisticated attacker.

WWAN security is important for service providers and consumers. Most enterprises do not rely on native WWAN security since their requirements mandate end-to-end security. Instead they overlay the network with a VPN or SSL connection that provides the needed authentication and encryption and thereby shields them from any vulnerabilities In the air interface.

6.5.2.5. Unauthorized Wireless Bridging

The problem of unauthorized wireless bridging (illustrated in Figure 4-7) is simple. A PC with two network interfaces connects one interface to the private LAN and the other interface to a public network. The fact that the device is simultaneously connected to both networks is not inherently a security risk. However, it does become one if the user bridges the connection, for example through Internet Connection Sharing (ICS).

In practical terms bridging might be carried out by a laptop user with a Wireless Wide Area Network card (e.g. GPRS or 1xRTT). The user might have legitimate reasons for sharing the WWAN connection to the LAN, for example, for home use. However, in the office, the operation will create a disruption at minimum and can potentially provide unlimited network access to hostile intruders.

Figure 4-7
Unauthorized wireless bridge

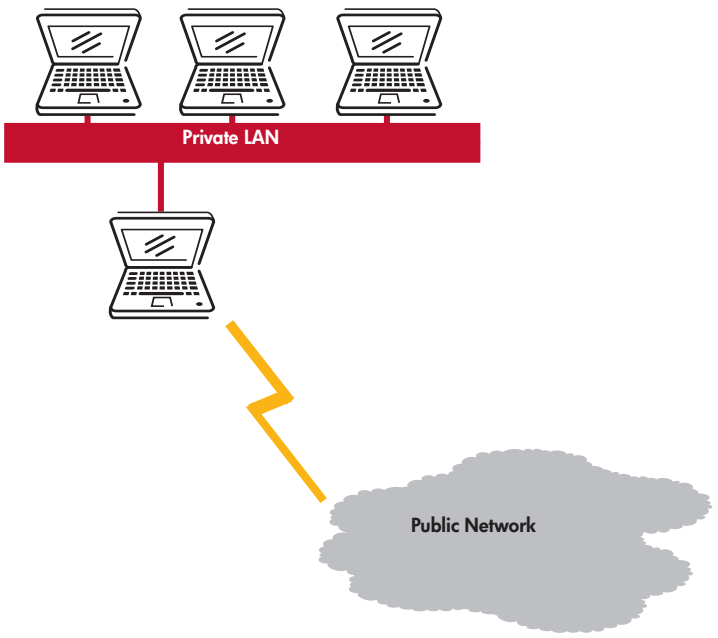

6.6. IPv6 Security

The Internet Protocol version 4 (IPv4) is the foundation of the vast majority of today's networks. Any change in this base protocol has serious consequences for all computerized applications and infrastructure.

The initial proposal to overhaul IPv4 was based on the expectation that IPv4 addresses would soon be exhausted. While that specific risk was temporarily deferred through Network Address Translation (NAT), IPv4's successor IPv6 has differentiated itself through some additional technical advantages that provide a powerful foundation for the creation of new and improved net-centric products and services.

- Scalability – larger address space
- Mobility – seamless roaming
- Administration – auto-configuration for network resources, renumbering of addresses
- Robustness – extensibility, and more

The need for new functionality will probably drive only a limited number of our companies to IPv6 in the short term. Some users will find they have no choice as they will find it increasingly difficult to obtain additional IPv4 address space. For more information, see the warning from the American Registry of Internet Numbers (ARIN) at **www.arin.net/announcements/20070521.html**.

Others may need specific IPv6 enhancements for specialized applications like mobility. For most enterprises the return on investment at this point is not yet compelling. They are, however, increasingly concerned about security. IPv6 is enabled by default with many Linux distributions as well as with Microsoft Windows Vista and Microsoft Windows Server 2008, which means that most enterprises are able to deploy some IPv6 technology today. If not carefully monitored and managed, rogue deployments can lead to devastating security breaches. The only way to address these threats is to develop a comprehensive security architecture that includes proactively monitoring any IPv6 network activity.

Additional developments that indicate the imminent adoption of IPv6 across the industry include:

• A U.S. White House directive (**www.whitehouse.gov/omb/memoranda/fy2005/m05-22.pdf**) states that "all agency infrastructures (network backbones) must be using IPv6 and agency networks must interface with this infrastructure" by June 2008. Outside the U.S., many government agencies across Europe and Asia now have similar directives which are forcing a ramp-up to IPv6.
• Mobile phone operators are deploying new technologies like third generation data networks or the IP Multimedia Subsystem (IMS) which require IPv6 in the core network.
• The European Space Agency (ESA) has declared its support of IPv6.
• The Japanese Intelligent Transport System (ITS) project supports IPv6 and the European Car2Car consortium has recommended exclusive use of IPv6 for future Car2car applications.
• The Digital Video Broadcasting (DVB-S) consortium has decided to move to IPv6.
• The Chinese Government has created and financially supports CNGI, an IPv6 backbone network designed to be the core of China's Internet infrastructure.
• CENELEC has opted for IPv6 for the Smart home concept.

As we prepare for the eventual deployment of IPv6, we must ensure this transition receives a systematic security assessment.

Fortunately, security has always been one of the highest priorities of the IPv6 architects. Its protocols are regularly scrutinized for any oversight that could jeopardize a wide-scale deployment, whether from risks related to confidentiality, integrity or even availability.

However, this is not sufficient. Every new technology brings with it new usage models, product portfolios and administrative requirements. The onus is therefore on every security architect to understand the security ramifications of emerging technologies and develop a plan that applies existing policies to the new requirements, or indeed updates the policies to reflect new demands and usage models.

6.6.1. Network Address Translation

One of the primary benefits of IPv6 is an enormous address space that can potentially remove the requirement for any Network Address Translation (NAT). The original purpose of NAT was to permit address growth beyond the range which a company or service provider was able to acquire an allocation. This requirement grew acute in the 1990s when it became apparent that the IPv4 address space would not be able to cope with the full set of network nodes that would eventually connect to the Internet.

In the meantime, many have also come to see NAT as a form of security, since it provides a control mechanism for incoming data, hides internal topologies and prevents direct connectivity between internal and external systems. In a small or home office environment a NAT router may act as a simple firewall, and in fact, NAT and firewall functionality are often combined on consumer gateways.

If IPv6 renders NAT unnecessary it would be easy to conclude that IPv6 weakens security. However, there are some logical counterarguments that address this concern:

- IPv6 removes the requirement for NAT and the IPv6 community discourages its use but it is still technically possible to use NAT on IPv6. IPv6 has no provision for universal private address space (RFC 1918) but there is no need since there is sufficient routable private address space.
- Any commercial firewall can enforce the same restrictions as a NAT appliance and can provide much more refined access controls in a flexible and secure manner. The Internet-Draft IPv6 Network Architecture Protection offers a complete comparison of the two approaches.
- Devices/appliances which perform NAT functions are often bundled with other functionality like firewalls and routing. The removal of NAT does not impact any of these functions.
- The focus in enterprise security is moving from the edge to other areas of the network. This does not mean that we can ever ignore the perimeter - it is still a critical ingredient of any multi-tiered defense. Nevertheless, comprehensive safeguards at other policy enforcement points allows more flexibility in balancing usability, cost, manageability and security at any given level.

In a compartmentalized network such as HP's Adaptive Network Architecture, or even a traditional DMZ, the notion of a single perimeter is weaker. There are many perimeters, and an attack from the hostile Internet will need to break through more than one of them to gain access to many assets. NAT adds substantial administrative complexity to such architectures, and the simplification provided by the unified, transparent IPv6 address space provides operational reliability benefits that outweigh the security benefits of NAT, as long as the "deny unless explicitly permitted" default firewall rule is enforced. NAT makes this rule technically difficult to defeat, which is a benefit in organizations with weak security management which may be swayed by business pressures to put expediency ahead of security.

6.6.2. Increased Address Space

One of the best-known benefits of IPv6 is its increased address length of 128 bits with the mathematical implication of an exponentially larger address space. This poses an insurmountable problem for many scanning tools. If the allocation is sparse and random then it is virtually impossible for a sequential or random probe of addresses to return a useful number of hits.

While this obstacle to scanning is generally a positive development it is important to note that it can also provide a challenge for legitimate tools. System and network administrators may currently work with tools that scan the network for inventory purposes, detecting malicious activity and even patch systems.

They must also resort to other techniques but their task is considerably easier than that of an intruder. Authorized administrators can obtain valid IP addresses from many network components such as routers and DHCP servers. Nonetheless this is an item that should not be ignored as the planning for IPv6 begins.

6.6.3. Hacker Tools

As IPv6 is coming into the mainstream, nobody would expect the hackers to be far behind. But perhaps they are actually ahead of most security departments. As described in "Security Implications of IPv6" by Michael Warfield, underground sites offer many IPv6-enhanced versions of old tools, such as halfscan6, netcat6, NMAP, Ethereal, Snort and TCPDump. They also host a number of protocol bouncers such as relay6, 6tunnel, nt6tunnel, and asybo that can relay and redirect connections between the two protocols, obscuring traceability.

6to4DDos is a distributed denial-of-service attack that leverages 6to4 tunneling to destabilize both IPv4 and IPv6 sites.

6.6.4. Transition

Transitions are always points of vulnerability. Multiple protocols are running on the network thereby increasing the exposure. IT has a double management load and a new technology to work with making the chance of human error greater. But there is no way to progress without transitions, so the best we can do is to identify the risks and try to address them.

One essential first step is to ensure that all of our current security policies and tools will continue to work in a transitional environment. The means that all our firewalls, proxies, intrusion detection systems, anti-virus scanners and network management software must be IPv6-compliant. Fortunately, most of the products in these categories have included IPv6 support. However, a prepared IT department must make a systematic check for the full feature set of each component.

The second major area of concern is around the transition mechanisms. There are many options, including 6to4, ISATAP, Teredo, configured tunnels, MPLS, GRE and DSTM. There isn't always unanimous agreement, because the security mechanisms are often one of the areas in which the drafts are changing most frequently.

Before considering even a pilot implementation, it is absolutely essential that the network security architect diligently assesses the options and ensure that they are functional and do not represent a risk in the particular environment. This may include checking firewall compatibility, authentication mechanisms, availability (redundancy and failover) and even the complexity of the solution.

Perhaps the most important consideration with regard to transition mechanisms is that they allow users to begin deployment of IPv6 ahead of their IT department. Almost anyone can easily begin to experiment with IPv6 by typing in a short line at the command prompt. In most cases, this activity will be confined to the local area network, and even though not centrally monitored, is fairly innocuous.

However, with many of the transition mechanisms being enabled by default, this danger is becoming more pronounced. It is now no longer a safe assumption that the connectivity is limited to the LAN. It is quite feasible that they will establish a connection to a public IPv6 network and potentially even draw legitimate users through their tunnel.

6.7. Best Practices for Secure Networks

6.7.1. Management

Management best practices for secure networks include well-defined and enforced policies, standards, user training, procedures, standard locked-down baseline configurations, and guidelines. Other areas of concern are extranet user agreements and the proper handling of worker termination. Specific engagements such as security reviews, risk and vulnerability assessments, and incident and event management must also conform to industry best practices. A good reference for security management practices is the "Information Security Management Handbook", by Harold F. Tipton and Micki Krause, Auerbach Publications, 5th Edition 2004, ISBN 0-8493-1997-8.

6.7.2. Operations Security

A trusted infrastructure depends on operational continuity and sustainability. This requires a consistent approach. Best practices are important for selecting and deploying infrastructure components and ensuring sufficient capacity to establish a robust, scalable, and highly available infrastructure. Operational sustainability also provides common supporting operations for backup, disaster recovery, replication, and business continuity.

Operations security represents the controls and safeguards that secure an enterprise's information assets on a computer or linked with the computer environment. Security controls address software, hardware, and processes. As the core component of information security, operations security controls the way data is accessed and processed, and it represents a set of controls designed to provide effective levels of security.

Operations security provides consistency across all applications and processes. It includes protection of physical assets, such as computing equipment, networks, and media. Operations security also includes resource protection, accountability access and use, and audit trails.

• Resource protection prevents the loss or compromise of an enterprise's computing resources, including main storage, storage media, communications software and hardware, processing equipment, standalone computers, mobile devices (as appropriate), and printers. Resource protection helps reduce potential damage from unauthorized disclosure and alteration of data by limiting opportunities for misuse.
• Accountability access and use ensures access for a specific authenticated and authorized individual user or system at a particular moment in time, and it tracks access and use to that individual or system.
• Audit trails track activity to specific individuals or systems to determine accountability and responsibility.

Operations controls for protecting resources require accountability and responsibility for all of those involved in developing, maintaining, and utilizing processing resources.

6.7.3. Physical Security

Physical security consists of controlling access to physical assets such as buildings, computers, and documents. Such assets can hold sensitive information and provide access to networked resources. Enterprises implementing physical security must plan the appropriate level of security for and access to site locations, buildings, computer rooms, and data centers. Managing and monitoring these facilities is a major component of physical security. To address physical security needs, enterprises must define:

• Best practices for the management and monitoring of physical facilities
• Mechanisms for protecting removable media and offline data storage
• Methods of securely labeling and protecting confidential documents

6.7.4. Firewalls

Firewalls secure network perimeters, workgroups, and hosts. They can be configured to block unauthorized incoming and outgoing traffic, conceal system identities and network topologies, log traffic, and log events of interest. Some firewalls have routing capabilities to direct incoming traffic appropriately, and some firewalls are used to authenticate network users. However, firewalls cannot defend against attacks that do not use the network or that use it in an authorized fashion. For example, an internal attack or malicious code downloaded from the Internet.

Firewalls can operate on a variety of platforms, including general-purpose servers, dedicated appliances, and desktop computers. The OSs of general-purpose servers must be carefully hardened to provide a secure environment for the firewall. This hardening process generally involves setting system parameters and disabling unnecessary system services. This process is not necessary for appliance-based firewalls, which come with their own vendor-configured and supported hardware. Desktop firewalls control traffic to and from the host upon which they reside, and they are installed directly on the desktop computer.

There are several different types of firewalls, including packet filters, circuit-level gateways, stateful packet inspection firewalls, and application proxy servers. They use different techniques to determine whether traffic should be allowed or blocked, and they operate at different layers of the Open System Interconnection (OSI) standard reference model set forth by the International Organization for Standardization (ISO), as noted in Table 4-4. For details about these firewall types, refer to Appendix B.

Table 4-4
Types of firewalls and OSI layers of operation

Type	*Layer of operation (OSI model)*
Packet filters	Network
Circuit-level gateways	Session
Stateful packet inspection firewalls	Network, transport, potentially others
Application proxy servers	Application

6.7.5. Network Architecture and Compartmentalization

Network compartmentalization is a fundamental design principle - a best practice for architecting networks to protect assets in accordance with key business drivers. Compartments are not necessarily based on physical location. Rather, they are logical groupings of assets that require similar levels of protection, regardless of location. Because it enables centralized management, compartmentalization greatly improves staff resource efficiency and promotes rapid and effective responses to security incidents.

Compartmentalization is accomplished by assigning an IP address space to each logical compartment, providing technological and geographical transparency. Access Control Lists (ACLs) and firewall rules (collectively known as policies) can be applied broadly to the compartment, rather than managed individually across the enterprise. In some cases, the same security policy applies to multiple IP address spaces at geographically different locations. This is known as a policy domain.

In time, the IP address of each network device determines its membership in a specific policy domain. Network traffic from and to a device is treated according to its domain's policy. HP's ANA, which is covered in more detail later in this chapter, supports this approach. Appendix A contains a more detailed description of compartmentalization.

6.7.6. Authentication, Authorization, and Auditing (AAA) Servers

Triple-A (AAA) servers authenticate network users, authorize them to use particular network resources, and account for their network usage. They provide a central control point for external network access, and they work with various types of network access servers that interact with users and collect their credentials. AAA servers are mentioned here because network access is a key element of a trusted infrastructure. The Identity Management chapter (Chapter 3) of this handbook provides further details about AAA technology. See Appendix C for a summary of AAA protocols.

6.7.7. Network Access Control (NAC)

NAC is primarily a network security element, intended to protect the network and its resources from harmful users and systems/devices. NAC controls and restricts access to network resources based on certain criteria and business policies. In its most basic form, NAC allows a network administrator to restrict network access to authorized users and/or devices. However, many organizations have the need to provide, or can benefit from providing, different levels of access depending on the role of the user.

For example, employees have access to internal network resources and the Internet while guest users are only provided access to the external Internet. There is also a need for protection from malicious software, which is accomplished by evaluating the health or security posture of devices connecting to the network. The required posture is defined by organizational policies and is based on checking for things such as operating system version, patches, security software (anti-virus, anti-spam, firewalls, etc.), security settings on common software installations, or other required or prohibited software. NAC goals can be further complicated by the fact that today's network is often comprised of network access requests from devices that are not under direct organizational control, such as contractor and guest laptops. Furthermore, the need to understand and comply with regulatory agencies and company policies alike drives a need for the organization to seek solutions that meet this goal, often with fewer resources than ever before.

Critically, NAC is not an isolated security solution, but is part of a layered security, or Defense in Depth approach to protecting your organization's information technology assets. The challenge of effectively controlling access to the network as part of an overall network security architecture based on business environments and goals emphasizes the importance of HP's extensive and proven strengths in this area. The key is not simply to understand the various components in isolation, but to know how best to combine and apply them to meet business needs. A primary goal of network access control technology is providing a mechanism for deciding whether to allow or deny access based upon endpoint compliance with relevant configuration guidelines and business security policies.

Complete NAC solutions incorporate appropriate endpoint, edge, core, LAN and WAN controls. NAC also provides mechanisms to quarantine and remediate non-compliant devices to allow them appropriate access to network resources. A managed endpoint is any device that connects to the network, and is subject to the compliance of security policies and is under administrative control. An unmanaged endpoint is any device that connects to a network, is subject to the compliance of security policies but is not under an organization's direct control (e.g. a contractor computer, guest user, or other). An Endpoint Policy is a collection of tests, or criteria, used to evaluate the integrity of an endpoint device attempting to access the network.

There are two primary models for NAC implementations: pre-connect NAC and post-connect NAC. Implementations may utilize hybrids of these two.

- Pre-connect NAC refers to network access control where the testing of the device to ensure compliance with network access policies is done prior to the device being granted regular access on the network. A similar term used in the industry is "pre-admission".
- Post-connect NAC is network access control where validation and monitoring of the endpoint continues after access to the organizations network has been granted. A similar term used in the industry is "post-admission".

Although NAC is about controlling access to the organization's network, control should not stop after pre-connect checks. Even a healthy endpoint with a known and trusted user has the potential to cause consequences once on the network. This drives the need for a more complete NAC solution that provides post-connect NAC validation and monitoring.

Figure 4-8 depicts an overview of network access management components. When a device attempts to connect to the network, via a wired jack, a wireless AP, or a remote access service, a policy enforcement point (PEP) allows or denies access. The ultimate access decision is made by a policy decision point (PDP), which consults the relevant enterprise policies to make an informed and correct access decision. A policy repository houses the policies and rules consulted by the PDP. In addition, the process generally includes an endpoint assessment step to determine the health of the system attempting access. If the device attempting access is deemed unhealthy or unauthorized, it can be routed to a special quarantine area separate from the enterprise network.

Once a user or device is permitted to attach to the network, policies based on the user or device identity and/or role determine the level of access; for example, authorized vendor A is permitted access to systems x, y, and z within a defined time window.

Figure 4-8
Network access management: logical view

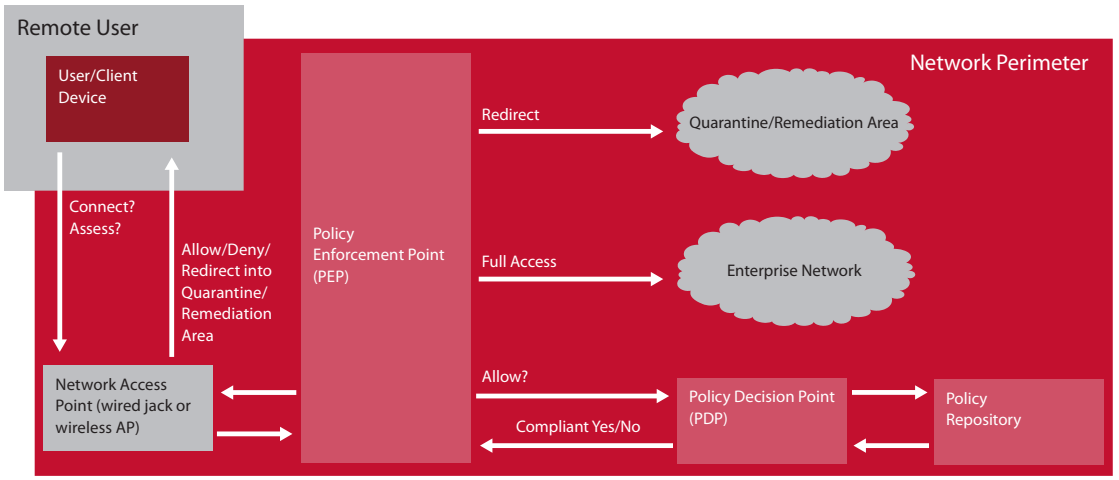

6.7.7.1. Policy Decision Point (PDP)

Enabling effective access policy decisions can require collecting a variety of data and policies that are spread among numerous systems across the enterprise. An important aspect to consider is the identity of the user or process generating the access request, for example:

- Is the user or process known or unknown to the enterprise?
- If known, what privileges are associated with this user identity or process?
- If unknown, what if anything can be deduced from the context or other attributes associated with this entity?

Even when the user associated with the device is known (successfully authenticated) and authorized to access the entire network or specific hosts or services, a separate evaluation of the device is necessary. Evaluation can include various attributes of the device, determined through endpoint assessment, such as:

• Device type: Computer, OS, PDA, any VOIP device, smartphone, or other?
• Registered or known device: Is the device known from past encounters or does it have a registered device identity?
• Resource access policy: If authentication of the device is possible, does the use of this device meet the policy for access to the requested resource?
• OS and application updates: Are the following items verifiably up to date?
    • OS and application patches
    • Anti-virus software and definitions
    • Anti-spyware software and definitions
    • Device (personal) firewall

6.7.7.2. Policy Enforcement Point (PEP)

Once the PDP reaches an access decision, the resource to which the device is attempting to connect enforces the decision at the network level. If a device and user are known and authorized to connect to the network, the PEP checks for restrictions on the access. Restrictions can include:

• Access valid until a certain time/date
• Access limited to certain network zones or compartments
• Access limited to specific resources

6.7.7.3. NAC Deployment Options

A good NAC implementation will utilize pre-authorization checks against security policy in order to protect the network from harmful systems. The following enforcement modes are common methods for deploying NAC and can be used together to provide complete access control coverage across the network:

• **Endpoint Enforcement**: Using endpoint enforcement entails having an agent deployed on the endpoint making decisions based on policies it has been given, before the device or user is allowed to connect. While appealing from a management standpoint, this is not necessarily true NAC as the environment must fully trust the endpoint anyway.

• **Edge Enforcement**: Using a combination of either 802.1x, SNMP or CLI to control switches, or through a VPN gateway, edge enforcement based NAC is often the most efficient and effective enforcement method and is recommended for environments with devices. Often 802.1x is an increasingly common method used to authenticate devices that are connecting to the network. In this mode, users and devices are usually but not always authenticated using RADIUS. Endpoints are isolated so they can be tested for security policies. Then, they are either allowed to join the network, or are put in a remediation network so the user can resolve the security settings that have caused the isolation.

• **DHCP Enforcement**: NAC solutions can integrate with an organization's DHCP servers to isolate and test endpoints. As endpoints request a network address, they are isolated by their network address so they can be tested for compliance with security policies. If they comply, they are provided with a new network address and allowed to participate on the network. If they fail, they are placed into a remediation network so the user can resolve the security settings that have caused the isolation. This method is useful for environments where 802.1x is not available because it is not supported by the network infrastructure.

• **Inline Enforcement**: In this mode, NAC solutions are placed inline with network traffic and actively filter new connections until they are tested for compliance with the security policies. This is an effective solution for testing endpoints that connect remotely through a VPN concentrator.

• **Out-of-band Enforcement**: In this mode, NAC solutions are placed deeper in the network either at concentration points or often in the data center itself. An out-of-band solution monitors traffic on switches to determine if enforcement is needed, and when required, will enforce policy by changing PEP (e.g. switches or firewalls) configurations. For example, some vendors integrate with existing deployments using HP's sFlow and Policy Based Routing (PBR) controls, or Cisco ACL's through CLI commands. Limitations are potential lags in detecting policy violations allowing time for infected endpoints to attack others, but deployment options are sometimes cheaper.

Additionally, some deployments use no enforcement at all, at least, initially. The goal here is to enable the logging and reporting aspects of the NAC solution, such that security and IT management can review, learn and determine the actual end-user impact of policy enforcement. In this way any unforeseen and potentially negative consequences can be avoided before they occur. Choosing to ratchet up the level of enforcement over time, in terms of what restrictions are enforced and what remediation is required is another way to carefully roll out NAC. For example a staged approach may require locking down the core finance department network immediately while only warning or alerting other departments about their lack of compliance without lock-out for a certain time.

Each mode provides some benefits and poses some drawbacks to the security of certain networks. For example, the inline mode has a greater ability to restrict devices since the PEP physically sits between the clean and unclean networks; however, this mode can sometimes to be hard to scale up to larger deployments depending on the vendor management tools and hardware deployment options. The DHCP model is well-suited for existing infrastructures of any size, but care and consideration must be given to the current network's threat model for this model to be effective. Utilizing IEEE 802.1x provides a robust authentication scheme that integrates well, but it requires extra infrastructure (such as RADIUS services and 802.1x supplicants).

Ultimately deployment options must meet the needs of the governance and security model of your organization. As a result a combination of these solutions may be required to meet those needs, in parallel (layered) as well as across different locales. For example, using both endpoint and in-line enforcement allows for flexibility when dealing with areas that allow for both corporate owned and customers devices connecting.

6.7.7.4. Quarantine and Remediation

In order to support business functions, a network cannot simply deny access with absolute quarantines; remediation mechanisms must also be included in NAC to facilitate the remediation of endpoints. If a device attempts to connect to the network, and it does not meet policy-based criteria for permitting access, it can be isolated to a quarantined area (or compartment) of the network that allows access to certain unprotected resources. In some situations this may be sufficient, as in the case of guests attempting to simply connect to the Internet to check external e-mail accounts or access public web sites. Some organizations provision a separate network segment dedicated to guest access, with no connection to the internal enterprise network.

If, however, the user is authorized and desires access to protected resources, but problems with the device have been identified through the endpoint assessment, the quarantine area could provide access to resources for correcting identified deficiencies. This might include access to current anti-virus or spyware definitions, OS or application patches, and device-specific or personal firewall updates. The device is allowed on the enterprise network only when the identified deficiencies are addressed, with access determined by the policy affecting the user or process.

Some organizations may decide that entry into the quarantine area is not automatic, due to the possibility of the quarantine becoming a source of infection. An additional concern is acceptable use policy (AUP) issues, such as the organization's open connections becoming the source of inappropriate (or even harmful or illegal) traffic.

6.7.7.5. Assessment Methods

When attempting to connect to the network, an endpoint must be assessed to determine whether it is in compliance with relevant security policies. There are two fundamental approaches to assessing the status of a device attempting to access the network, active or passive.

With active assessment, a software agent is installed on the device. It performs certain tests, gathers information, and reports its findings back to a server dedicated to processing the information to determine the most effective means of remediating any problems identified. This method can be problematic for a number of reasons, including resistance to installing software unknown to the user on the system. Given the challenges associated with supporting software on potentially unknown and unsupported systems, this approach may not be viable for all devices attempting access, particularly devices such as PDAs or smartphones. There are three primary ways in which active assessment occurs:

- Agent-based Permanent: This requires software to be installed on each endpoint and once installed and running it is always available for the endpoint to be tested.
- Agent-based Transient: This requires an agent to be temporarily downloaded to the endpoint while it is being tested.
- Agent-less: The agent-less approach uses native applications or APIs to provide agent functions that are then used for testing from a PEP.

Passive assessment is performed without installing software on the device, but rather by assessing its responses to certain stimuli from the network. Requiring devices to install software agents to determine and report on their health or install patches or updates is potentially problematic. The requirements may violate guests' home organization policies, and the definition of system health can vary from organization to organization. Ultimately, network operators must determine the appropriate policies required for systems to attach to their networks. Users must decide whether the policies are acceptable and whether they are willing to connect under the terms required.

6.7.8. Intrusion Detection and Prevention Systems

Networks are still vulnerable to external and internal attack, even if they are properly secured and every effort is made to control host security. IDSs and IPSs provide an extra layer of defense. An IDS detects and reports exploitation of network and system vulnerabilities, whereas an IPS detects such exploits and takes immediate action to thwart them.

IDSs may be host-based or network-based. Host-based IDSs reside on servers and analyze audit logs and other indicators of system activity. Network-based IDSs use dedicated hosts that intercept and analyze network traffic. IDSs detect intrusions and other exploits such as privilege abuse by using predefined rules, predefined attack signatures, or observed deviations from normal activity (statistical anomalies).

IPSs take the idea of an IDS to the next step. After detecting an attack, an IPS performs specific actions to block an attempted attack or render it worthless. Like an IDS, IPSs may be host- or network-based. IPSs may respond to an attack by dropping suspicious data packets, terminating suspicious sessions, denying user access to resources, reporting activity to other hosts that may also be vulnerable, or updating their own configurations to better address specific attacks. IPSs can integrate with firewalls, so that when an IPS detects a source of hostile traffic, the firewall works to block it.

6.8. HP Network Security Products and Solutions

This section provides an overview of HP's offerings, services, and solutions related to network security. HP provides a broad set of product and service offerings in this space, including many supplied by trusted HP partners, in order to bring the best possible solutions to our customers. After facing many of these challenges across our own internal networks and working with a number of diverse customers large and small, HP is uniquely positioned to bring considerable expertise to the enterprise.
For further information, please see
www.hp.com/go/security/trusted.

6.8.1. Adaptive Network Architecture

HP's ANA is a blueprint for:

- Segmenting or logically compartmentalizing an enterprise network based on the business needs of applications or hosted services
- Extending the compartments enterprise-wide as required regardless of physical location
- Enabling centralized policy management for the resulting architecture

Conventional perimeter defenses can no longer strike a balance between fast-changing business needs and sufficient protection of company information assets. Today, management of enterprise firewalls is typically exception-based, with a large number of access holes that accommodate specific user or system requirements. These exceptions cause both security and operational concerns. ANA transforms legacy enterprise data network architecture from a monolithic perimeter to a set of purpose-built (and more secure and manageable) distributed network compartments.

Compartmentalizing is not new; enterprises have been doing it for years but in a limited fashion. Traditionally, it has not been cost effective for companies to compartmentalize the entire network-conventional approaches are not scalable or sustainable. As a result, companies only compartmentalize a small portion of their network. Implemented internally within HP since 1999, ANA breaks through this barrier by combining processes, technology, and a governance model. The governance model is a tested, hierarchical arrangement, where business units, IT architecture, network engineering, and network operations interact at various levels to instill agility and consistency for planning and executing change to network access policy.

There are three areas that demonstrate the capability of ANA: IP Communications (IPC) deployments, network consolidation (In the context of data center consolidation), and network admission control (802.1x) adoption. In all cases, a generic access policy must be applied hundreds of times throughout the network. As business needs evolve and change, modifying such a policy on a global scale is arduous. ANA enables agility by providing a means to manage policy centrally and enforce it in a decentralized way. With ANA, changing a hundred geographically distributed networks to permit or deny a specific application service can be handled in hours instead of the days or weeks needed by conventional practices.

HP has filed several patents for the process and design elements that form the underpinnings of ANA and has successfully deployed ANA worldwide for internal operations. ANA has enabled HP to reduce network administration costs and operating expenses, while shortening lead times for acquisitions and external collaboration.

Those interested in ANA have two implementation options. HP Services can provide the design, planning and implementation of ANA in the context of network solutions such as IP Communication, Network Consolidation and Network Security. HP Services will also deliver ANA as part of a complete outsourcing solution. Outsourcing Services delivers ANA to outsourcing customers.

6.8.2. HP ProCurve Networking

For nearly 20 years, HP ProCurve Networking has built enterprise LAN products that help people run their businesses more effectively. By providing a complete and affordable portfolio of network security solutions and services, alongside HP's highly skilled professionals, these products can help manage information resources, provide consistent performance, and deliver secure access to the enterprise. HP ProCurve's networking products support the ProCurve Adaptive Edge Architecture (AEA), which moves intelligence and security to the edge of the network where users connect. Enforcing security at a central point gives malicious traffic an opportunity to infiltrate the network from the edge to the core.

Stopping any unauthorized or suspicious activity at the edge or access point immediately isolates the problem and reduces the chance that the network as a whole will be impacted. This approach prevents users from gaining unauthorized network knowledge or performing electronic snooping to uncover passwords or other critical information that might assist in a network attack.

HP ProCurve's unified approach addresses both wired and wireless access and secures end-user access methods to the enterprise LAN. If a security breach cannot infiltrate the host because network intelligence locks out the potential attacker, enterprise network security improves dramatically.

6.8.2.1. HP ProCurve Networking: ProActive Defense

HP ProCurve's security framework is called ProActive Defense. ProActive Defense is the HP ProCurve approach to delivering a trusted network infrastructure immune to threats, controllable for appropriate use, and able to protect data and integrity for all uses. HP ProCurve delivers security solutions that proactively prevent security breaches by providing comprehensive access control (access security) and integrity and confidentiality of sensitive data (privacy).

HP ProCurve delivers security solutions that defend the network by securing the network from unauthorized extension (infrastructure security), by keeping the network safe from virus attacks (network immunity), and by providing reports to administrators with valuable information about the security of the network and security policy compliance (command from the center).

Access security controls which users have access to systems and how they connect in a wired and wireless world. HP ProCurve provides:

- Standard 802.1x port-based access control for all HP ProCurve enterprise-class managed products
- A combination of 802.1x with 802.1Q standards for two levels of security - when a user authenticates via 802.1x, HP ProCurve switches can easily place the user on the appropriate virtual local area network (VLAN) based on information from the authentication server, which limits users to exactly the network resources they are allowed to access
- 802.1x for 802.11 wireless networks to ensure only authorized users are granted access to the enterprise network
- Restrictions efficiently implemented at the network edge to control access rights and privileges that each user or group has to specific network resources, such as individual subnets, servers, or applications

Privacy ensures the integrity and confidentiality of sensitive data. HP ProCurve provides:

- Protection from data manipulation
- Prevention of data eavesdropping
- End-to-end VPN support for remote access or site-to-site privacy
- Wireless data privacy using 802.11i WiFi Protected Access (WPA), Wired-Equivalent Privacy (WEP), and VPN technologies

Infrastructure security includes protection of network components and prevents unauthorized users from overriding other security provisions. HP ProCurve provides:

• Secure controlled access to the configuration and management of the network infrastructure
• Switches that can authenticate network managers in a number of ways
• Protection of remote management access to the console prompt using the Secure Shell (SSH) protocol

Network immunity relates to designing a network infrastructure to survive an attack without interrupting service. Network-based viruses can infect authenticated notebooks and PCs when they connect to the Internet outside of the office. In addition to attacking network components, these viruses can compromise the network from within. With HP ProCurve:

• Products come with a number of built-in features that improve the network resiliency in the face of virus outbreaks.
• Management functions are protected from broadcast storms, flooded traffic, and network loops, enabling access to switch management in the presence of these network anomalies.
• Products help reduce excessive broadcast traffic that impacts every station on the network and typically results from an erroneous situation.
• Software releases run through extensive testing before distribution - one of the many standard regression test suites includes the CERT Coordination Center (CERT/CC) vulnerability test suite that bombards a switch with network attacks.
• Routing switches support authenticated updates from authenticated routers.

Command from the center is the ability that HP ProCurve management tools provide to set security policies, report alerts, and report information about the security of the network.

HP ProCurve security solutions move important access decisions and policy enforcement to the edge of the network where users and applications connect. Core resources are freed to provide the high bandwidth interconnect functions they are meant for, which means enterprise networks are optimized to perform better. What is more, effective control to the edge helps enforce security policies necessary for network convergence and a mobile workforce.

HP ProCurve Networking solutions have several layers of built-in security and take advantage of the latest standards-based security features to protect data. HP ProCurve's diverse array of security products and services bring trust, reliability, and flexibility to enterprise networking.

6.8.2.2. HP ProCurve Networking: ProCurve Identity Driven Manager

Network intrusions do not always come from Internet connectivity. Because intrusions can originate from within the local network, network access controls are critical to network security. ProCurve Identity Driven Manager (IDM) builds on the secure access features of ProCurve hardware and the standard RADIUS authentication process.

Figure 4-9
HP ProCurve Identity Driven Manager 2.0

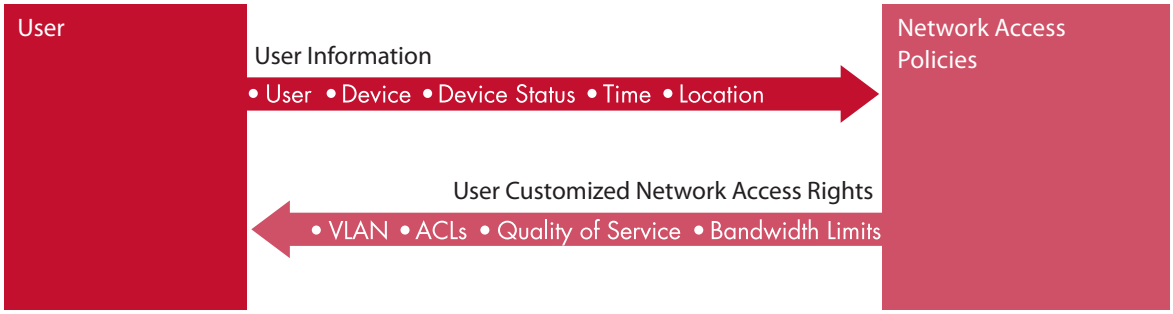

As Figure 4-9 illustrates, IDM allows businesses to define network access policies that enforce secure access to the network and provide dynamic security and performance configuration to network ports as users connect. These policies can allow, deny, and customize network access based on user, device, time, and location. The IDM 2.0 release adds integration with the TNC architecture. This integration allows network policies to verify that a system is compliant with business policies before the system is allowed into the network.

ProCurve IDM is a key piece of ProCurve AEA, providing the ability to centrally manage network access policies while controlling a secure, adaptive edge network.

6.8.2.3. HP ProCurve Networking: HP Virus Throttle Software

As every IT manager knows, computer virus epidemics are only getting worse. Current methods to stop the propagation of malicious agents rely on the use of signature recognition to prevent hosts from being infected. While this approach has been effective in protecting systems, it has several limitations that decrease its effectiveness as the number of viruses increases. Signature recognition is fundamentally a reactive and case-by-case approach. The latency between the introduction of a new virus or worm into a network and the implementation and distribution of a signature-based patch can be significant. Within this period, a network can be crippled by the abnormally high rate of traffic generated by infected hosts.

Virus throttling, in contrast, is based on the behavior of malicious code and how it differs from unaffected code. Normally, a computer makes fairly few outgoing connections to new computers and is more likely to regularly connect to the same set of computers. This is in contrast to the fundamental behavior of a rapidly spreading worm, which attempts many outgoing connections to new computers. For example, computers normally make approximately one connection per second; the SQL Slammer virus tries to infect more than 800 computers per second.

HP Virus Throttle software establishes a rate limit on connections to new computers. Normal traffic remains unaffected, but suspect traffic that attempts to spread faster than the allowed rate is slowed. This creates large backlogs of connection requests that can be easily detected. Once the virus is slowed and detected, technicians and system administrators have time to isolate and remove the threat.

A virus-throttle approach differs from signature-and-patch approaches in three key ways:

- It focuses on the network behavior of the virus and prevents certain types of behavior, in particular, the attempted creation of a large number of outgoing connections per second.
- It restricts the code from leaving the system instead of stopping viruses from entering the system.
- It makes the system robust and tolerant to false positives by allowing connections beyond the permitted rate to be blocked for configurable periods of time.

Figure 4-10
HP Virus Throttle process

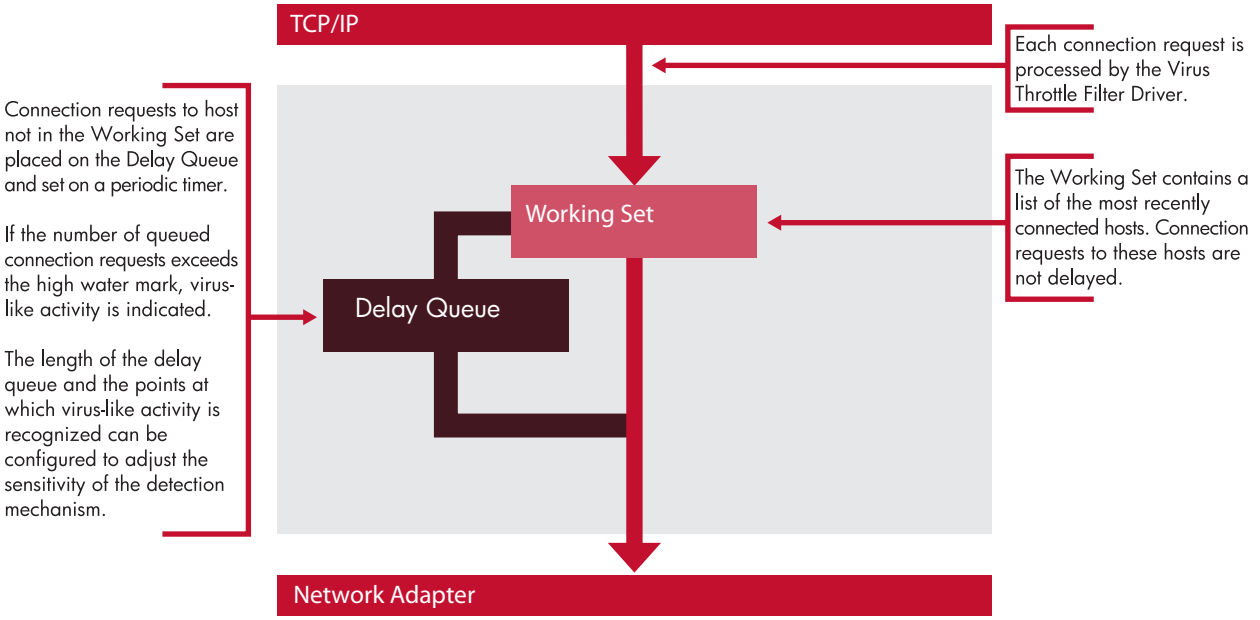

HP Virus Throttle software should complement, not replace, signature-based solutions. The virus-throttle technology fills a gap in anti-virus protection that has, until now, allowed previously unknown threats to wreak significant damage before patches can be deployed. With HP Virus Throttle, previously unknown threats can be mitigated, giving administrators time to deploy signature updates and patches. Figure 4-10 illustrates the process employed by HP Virus Throttle software.

6.8.3. HP ProLiant Essentials Intelligent Networking Pack

HP ProLiant Essentials Intelligent Networking Pack is a software solution available for HP ProLiant servers running Microsoft Windows 2000 and Microsoft Windows 2003. It offers advanced networking and combines capabilities for redundancy and load balancing. HP Virus Throttle software, described previously, is also implemented in this product. When implemented with other virus-prevention tools, HP ProLiant Essentials Intelligent Networking Pack provides an extra layer of protection against attacks that can bring down the entire network.

6.8.4. HP ProLiant DL320 Firewall/VPN/Cache Server

The HP ProLiant DL320 Firewall/VPN/Cache Server running Microsoft Internet Security and Acceleration Server 2004 provides an affordable, integrated, easy-to-use, and manageable hardware security and caching solution. It can be quickly deployed to help protect key business applications, such as Microsoft Exchange Server, Outlook Web Access, Internet Information Services, and SharePoint Portal Server. In addition, Microsoft Internet Security and Acceleration (ISA) Server 2004 integration with Windows Active Directory services enables administrators to use the solution to apply group- and user-level policy and authentication across a broad range of scenarios, including firewall policy, VPN authentication, and outbound web proxy and access control.

6.8.5. HP IPFilter

HP IPFilter filters IP packets that access HP-UX servers. IP packets are granted or denied access to or from the system based on stateful packet inspection and sophisticated packet-filtering rules. Featuring a unique packet throttling technology called Dynamic Connection Allocation, IPFilter can be configured to either prevent or minimize the effects of many types of DoS attacks. A HP Service Professional can remotely install and configure IPFilter on a qualified HP Integrity Server and verify that the software starts up and shuts down without error.

Table 4-5
Examples of HP partner secure network offerings

Partner	Product Name	Purpose	Partner Website
Cisco Systems	Cisco Clean Access	Enforces network security policies	www.cisco.com
	Cisco Secure Access Control Server (ACS) for Windows	Manages user access to Cisco devices and applications with 802.1x access control	
	Cisco VPN-Enabled/Optimized Routers	Supports IPsec VPN features within Cisco routers	
	Cisco PIX 500 Series Firewalls	Provides stateful packet inspection, IPsec VPN, IPS, and other solutions for a wide range of device applications	
	Proventia and RealSecure product lines for IDS/IPS	Provides IDS and IPS solutions	www.iss.net
Microsoft	Internet Security and Acceleration (ISA) Server	Provides application-layer firewall capabilities, VPN, and web caching Integrated with HP ProLiant DL 320 Firewall VPN Cache server	www.microsoft.com www.hp.com/go/proliant
Nokia	Nokia Firewall/VPN appliances	Provides an integrated solution for secure Internet communications and access control using Check Point firewall and VPN software	www.nokia.com
Symantec	Symantec Enterprise Security Manager	Manages and reports on security policy compliance	www.symantec.com
	Symantec Enterprise Firewall	Provides an enterprise-level firewall for Windows and Solaris platforms	
	Symantec Gateway Security	Integrates stateful packet inspection firewall, anti-virus, IDS/IPS, content filtering, IPsec VPN, and hardware-assisted encryption technologies in a self-contained device	

6.9. HP Partner Secure Network Offerings

To provide a complete and integrated set of options, HP has partnered with leading vendors whose products and services enhance and complement HP products and services. Table 4-5 on the previous page summarizes some of these partner offerings. See **www.hp.com/go/security/strategy** for details and updates about HP partner information.

6.10. Network Security Summary

HP's approach to network security begins with rigorous, widely accepted analysis and planning techniques. Network design is based on proven, integrated solutions and leading products. For organizations that must adapt rapidly, HP's ANA technology secures key solutions within enterprise networks while enabling them to change quickly in response to business imperatives.

HP and its partners offer a broad range of security expertise, products, solutions, and services to help ensure that organizations are not damaged by a disruption or compromise of their information flow. For more information, see **www.hp.com/go/security/trusted**.

7. Storage Security

In principle, storage security is straightforward. In practice, establishing storage security requires specialized knowledge, careful attention to detail, and ongoing review to ensure that storage solutions continue to meet an organization's evolving needs.

7.1. Environment

Storage security represents a major component of the overall security plan for a data center and a business. Consequently, business policies and practices must augment any hardware- or software-level security model, including network and system security.

7.1.1. Threats

Storage has evolved into a resource shared by many systems on a network. In many cases, it is no longer sufficient to secure just one system to which a storage device connects, because storage devices now connect to many systems. To protect against a variety of threats (not all of which can be anticipated in advance); storage security must address the varying security requirements of a diverse number of databases and applications. For example, storage security must protect:

- Valuable data belonging to each system against unauthorized access, modification, or destruction by any of the other systems
- Storage devices themselves against unauthorized configuration changes, with audit trails of all such changes

There is no value in carefully securing storage and subsequently leaving the system wide open to the Internet. Storage security must be a part of an overall security plan, both for a single data center and for the organization as a whole. Storage security also consists of a set of procedures that define access rights for data and authority for managing devices, and it defines an appropriate response when security issues occur.

7.1.2. Types of Storage

There are three main types of storage to consider today and two emerging technologies:

- Direct Attached Storage (DAS) is connected directly to a single system, similar to the disk within a PC.
- Network Attached Storage (NAS) is accessed via the Ethernet LAN network, and it stores and retrieves files.
- Storage Area Network (SAN) storage is accessed over a storage network, which today is typically based on Fibre Channel architecture, providing what looks like disk drives to systems.
- Internet SCSI (iSCSI) offers storage networking over IP networks. It is an emerging technology being utilized in small-to-medium environments where I/O throughput demands are relatively low. iSCSI is an important addition to SAN technology because it enables a SAN to be deployed in a LAN, WAN, or MAN (local-, wide-, or metropolitan-area network).
- Object storage is an emerging technology that combines aspects of SAN and NAS.

HP's storage security focus is on storage shared between many systems on a network, primarily SAN and secondarily NAS. Storage security is not a box added to a SAN as a firewall is added to a network. Security must be an attribute of every system, every switch, and every device in the SAN.

7.1.3. Benefits

Storage security provides protection from attacks and resulting exposures. Specifically, storage security:

- Protects data confidentiality
- Protects data integrity
- Protects data from destruction or loss

7.2. Principles of Risk Mitigation

Many ways exist to gain unauthorized access to data and to retrieve, alter, or destroy data. Examples of risks that may require mitigation include:

- Stealing disk(s) and backup tapes
- Copying disks
- Allowing an unauthorized system to access a disk array or tape library
- Wiretapping within a data center and between data centers
- Making unauthorized changes to permissions in the disk array or in the switch
- System mounting and initializing a volume it does not own as a result of a software defect
- Operator error or miscommunication

7.2.1. Mitigation Techniques

Mitigation of storage security risks involves identification and authentication, authorization, auditing, encryption and key management.

7.2.1.1. Identification and authentication

Identification and authentication techniques include:

- User logon identification and authentication via security mechanisms such as user name and password protection for authorization of administrative actions
- Audit trails (logs) to identify what was done and by whom, which deter deliberate misuse of authority and help recover from incorrect actions
- Timely revocation of an individual's identity or modification of authorization when responsibilities change or the individual leaves the organization
- Device identification and authentication through emerging technologies that ensure a device is permitted on the storage network (These technologies can also detect an "impostor" rogue device pretending to be a different device or system.)

7.2.1.2. Authorization

Authorization techniques include:

- Authorization for an individual to manage only specific devices or to limit access to many devices
- Verification by storage devices that an administrator who issued a command is authorized to do so, before performing the requested action
- Verification by disk arrays that the specific system that issued a read or write command has permission to do so for that Logical Unit Number (LUN), before performing the input/output (I/O) action (Through emerging technologies, a tape library controller can similarly verify permissions on I/Os to a tape library.)

7.2.1.3. Auditing

Auditing techniques include:

- Logging all administrative actions (changes) and any significant events (This is typically logged individually within devices, but logging software is the preferred method because it presents a single view and allows queries.)
- Extending the auditing mechanism over the entire storage network to track activities related to each element.

7.2.1.4. Encryption

Encryption techniques include:

- Encrypting data at rest on media such as a disk or a tape.
- Encrypting data in flight between data centers (and potentially within a data center) to protect against wiretapping.
- Encrypting data in use at the database or application level. This last item is not always considered a storage security item, but rather an application- or service-specific security feature.

Customers are always faced with the question of what is the best place to encrypt. Trade-offs exist for each of the three encryption locations listed above. Customers will typically face more management and deployment complexities the higher in the stack they want to provide encryption (for example for data in use instead of simply data at rest) - certainly when they attempt to deploy encryption across a large environment. On the other hand: the higher in the stack they encrypt data, the more protection can be provided. Encryption also represents an impact on performance.

Encryption is gaining more widespread use, driven by the increased incidents of lost customer data and the growing amount of regulations. HP considers that initially customers will want to deploy encryption at the storage device or media level because it is cheap and simple to manage. Over time, however, enterprise data center customers will consider deploying encryption in multiple places based on what is legally required, what solutions are available and the solutions' operational impact. Today for example customers selectively and surgically encrypt at the database and application level. Moving forward, HP expects them to deploy encryption at the storage device or media level, because in the future doing so will become cheap and simple to manage.
The increased data threats and the need for faster data encryption also drive standardization efforts to build hardware encryption into storage devices. For example, data copied between data centers is no longer protected from wiretapping by the physical security of the data centers.

The lack of physical security on cables outside the data center can be mitigated by passing the traffic through a dedicated encryption system before it leaves the sending data center. Similarly the encryption of backup tapes that are transported off site is of increasing interest to many organizations and has lead tape vendors to build hardware encryption capabilities into tape devices. HP also expects that as time goes on, encryption standards will be increasingly driven into SAN switches.

Dedicated encryption systems are available for both Fibre Channel and IP networking, with the latter called Internet Protocol Security (IPsec) gateways. Because of their cost and complexity, such installations are not common today. In the case of iSCSI, HP anticipates that IPsec will be built into future interfaces, making encryption more affordable and more ubiquitous than is currently possible with IPsec gateways.

7.2.1.5. Key Management

When dealing with encryption in enterprise environments key management is an area of particular concern: if you lose the keys you cannot decrypt the data. A primary concern here for example is that keys neither be lost nor exposed while a tape is still retained: the keys need to be accessible as long as the data lives, and for data at rest that could be decades.

Key management becomes unwieldy if multiple disparate solutions are deployed or the amount of encrypting devices grows across the enterprise. What's needed in terms of data center key management is the following:

- A scalable enterprise-wide infrastructure solution from a trusted vendor. HP expects that the volume of encrypted data will mushroom over time.
- Open, standards-based key management that integrates across the enterprise and across vendors. Interoperability is crucial.
- Reliable and highly available key storage. This means with an ample provision for encryption key redundancy and without taking a performance hit.
- Secure access and control meaning the establishment of key management policies ensuring that only those with the right permission can access the keys.
- Simplified and automated management is a must because protecting data is already complex and resource-intensive. Key management needs to be as simple as possible, even at the enterprise level.

The summary is that in the context of data and storage security, enterprises will need an effective way to centrally and automatically manage keys in a secure fashion, independent of where encryption occurs.

7.2.2. Data Access and Management Measures

Mitigation approaches generally fall into data access measures and management measures. Figure 4-11 categorizes these mitigation approaches by data security and management security.

Figure 4-11
Storage security model

Data access approaches include device authentication, device authorization, and encryption. Management techniques include individual authentication and authorization, with audit trails and logging, and key management. Some of the items in these categories are routinely used today. Others represent evolving, leading-edge mitigations.

7.3. Secure Storage Priorities

As businesses set secure storage priorities for the coming years, securing the management interfaces of all devices is usually the highest priority.

The key priorities for storage security include:

- Facilitating secure management of ports and interfaces on elements such as switches and arrays by:
  - Using strong passwords and changing default passwords
  - Disabling unused management ports on devices
  - Enabling firewall management of LAN interfaces to block widespread access
- Enabling LUN security (for example, Selective Storage Presentation and LUN Masking) if applicable
- Using encryption to meet regulatory compliance and specific business objectives

A plan for storage security must incorporate people and procedures as well as equipment. It must fit with the overall data center and business plans. This means evolving the plan as both the situation and technology allow, training users accordingly, following the plan, and testing it.

HP Storage Security Self- Assessment Tool

The HP Storage Security self- assessment tool has been designed specifically to help organizations understand how well their business is prepared for managing risk to sensitive data in their storage and backup environment, as well as complying with data privacy regulations. The tool has been developed by the HP Consulting and Integration Security and Risk Management Practice based on industry best practices and their experience gained through the design and deployment of hundreds of data protection engagements globally. Customers can access the tool from the following URL: www.hp.com/storage/securityassessment

Using the tool organizations must answer a set of questions and HP provides them with a personalized report that documents where their storage and backup security controls are appropriate, and where additional focus may be required. Recommended courses of action are designed explicitly to reduce the risk commensurate with their business, and to better comply with regulatory or industry mandates to which their business must adhere.

Customers can select any or all of the six key storage and backup security elements they would like analyzed: governance, compliance, operations, technology, encryption, and/or key management.

7.4. HP Secure Storage Solutions

7.4.1. HP StorageWorks LTO-4 Ultrium 1840 Tape Drive

The HP StorageWorks LTO-4 Ultrium 1840 is HP's first tape drive that is capable of storing up to 1.6TB per cartridge while providing hardware-based encryption. The LTO4 drives coupled with the HP Secure Key Manager (explained in the next section) can deliver a very secure data privacy solution for offsite data media such as tapes.

Below is a short overview of the LTO-4 Ultrium 1840's main product features:

- Supports AES 256-bit encryption
- Provides easy-to-enable encryption for a secure backup and helps prevent unauthorized accessed of tape media if cartridges are lost or stolen
- Includes a large-capacity, fast-performing tape drive with 1.6 TB compressed capacity and 240MB/sec compressed data transfer rate that is enhanced by HP's exclusive dynamic data-rate matching feature
- Provides enterprise-class reliability. The LTO-4 Ultrium 1840 is highly reliable, with an MTBF of 250,000 hours at 100% duty cycle, and includes HP One-Button Disaster Recovery (OBDR)
- Assures Investment protection. The LTO-4 Ultrium 1840 comes with a single-server version of HP StorageWorks Data Protector Express. It is read/write compatible with LTO-3 media and read-only compatible with LTO-2 media.

7.4.2. HP StorageWorks Secure Key Manager

The HP StorageWorks Secure Key Manager reduces an organization's risk of a costly data breach and reputation damage while improving regulatory compliance with a secure centralized encryption key management solution for the HP LTO4 enterprise tape libraries (see also previous section). The Secure Key Manager automates key generation and management based on security policies for multiple libraries. This occurs transparent to ISV backup applications.

The Secure Key Manager is a hardened server appliance delivering secure identity-based access, administration and logging with strong auditable security designed to meet the rigorous FIPS 140-2 security standards. Additionally, the Secure Key Manager provides reliable lifetime key archival with automatic multi-site key replication, high availability clustering and failover capabilities.

The HP StorageWorks Secure Key Manager provides centralized key management for HP StorageWorks Enterprise Storage Libraries (ESL) E-Series Tape Libraries and HP StorageWorks Enterprise Modular Library (EML) E-Series Tape Libraries.

In addition to the clustering capability, the Secure Key Manager provides comprehensive backup and restore functionality for keys, as well as redundant device components and active alerts. The Secure Key Manager supports policy granularity ranging from a key per library partition to a key per tape cartridge while featuring an open extensible architecture for emerging standards and allowing additional client types in the future needing key management services. These clients may include other storage devices, switches, operating systems and applications.

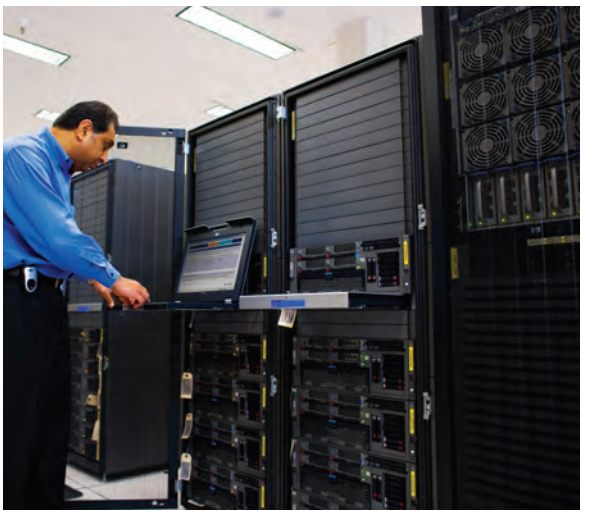

7.4.3. HP Compliance Log Warehouse

The HP Compliance Log Warehouse (CLW) can aggregate security logs enterprise-wide and tune security log reporting for specific audit and regulatory requirements. By understanding the detailed event data that IT systems already produce, organizations can better manage, investigate, and protect these systems. HP CLW collects and analyzes data such as system and application log files, database event records, and operating system event logs. With powerful compliance reporting tools, it turns this data into actionable intelligence, providing rapid time-to-value at a fraction of the cost of traditional data warehousing and security solutions.

The HP CLW is a high-performance appliance with log and analysis and real-time alert modules. It can provide high speed collection and analysis of log data that automates compliance reporting of many industry and government standards including SOX, PCI, FFIEC, HIPAA, NISPOM, DCID6/3, FISMA, EU Data Retention and ISO 17799. The HP Compliance Log Warehouse also has an adaptor for the HP Secure Key Manager.

7.5. Storage Security Summary

Storage security is part of HP's Trusted Infrastructure and is a major component of the overall security solution. Storage plays an indirect but critical role in an enterprise's overall security operations. A data center contains the majority of an organization's records; many business processes are affected if storage systems become unavailable or are compromised. An organization's storage and storage security strategy must relate directly to the business processes, IT infrastructure, and overall security model of the organization. Storage security draws not just on the organization's security governance and attitude toward risk, which is driven from a business level, but also on its centralized identity (authentication) and authorization services and its security management capabilities for managing threats.

In addition to mitigating security risks through independent identification, authorization, auditing, and encryption techniques tied to storage, a broader plan for infrastructure security across storage, networking, and hosts must also be in place. An attacker will seek weaknesses across all three areas. Securing storage over standard networking depends on how effectively the network is protected and on the security of the storage system itself. This is particularly true when storage is accessed over the organization's backbone network rather than through an isolated storage network or subnet.

8. Imaging and Printing Security

Security of the imaging and printing environment has long been ignored by IT administrators. Printers and scanners have been considered little more than network appliances, posing none of the risks of client and server PCs. Recent publications by hacker groups have raised the awareness that imaging and printing devices are more than simple appliances, and that these devices have capabilities beyond printing and scanning. As network printers and Multi-Function Printers (MFPs) grow in capability, they begin to resemble networked PCs in their ability to send and receive data. It would be wise for a company to view these networking devices like publicly available PCs with access to their network for sending and receiving data.

This section explains the threats and risks unique to imaging and printing environments and provides recommendations and strategies to prevent their effects. Parallels to common security capabilities are drawn to aid in explaining hardcopy-specific needs. Imaging and printing devices are put into the context of regulatory requirements, although - as will be seen - there is no simple solution.

8.1. HP's Imaging and Printing Security Framework

To simplify the presentation of security concepts, HP developed an imaging and printing security framework with three categories of security functions:

Table 4-6
HP imaging and printing security framework: security function categories

Secure the Device	Includes elements that protect the function of the physical device, including access controls for management and use, secure deletion of files, and physical security.
Protect Information on the Network	Includes network communications, including media access protocols such as 802.1x and secure management, scanning, and printing protocols.
Effectively Monitor and Manage	Includes the capabilities to securely manage fleets of imaging and printing devices and audit devices for compliance to security policies and regulatory requirements

The categories within HP's imaging and printing security framework are built from traditional network security theory, which identifies the four elements that compose a secure system: confidentiality, access control, integrity, and non-repudiation.

8.2. Secure the Imaging and Printing Device

Secure the Imaging and Printing Device includes capabilities that provide access controls to the functions of the device and ensure the integrity of its operations. Access controls limit MFP and printer functions to authorized users and include:

- Walk-up capabilities such as copying and digital sending
- Network printing
- Physical access to printed documents

Authentication requirements vary by environment, as do integration requirements to existing authorization mechanisms.

8.2.1. Multi-Function Printer (MFP) Walk-up Authentication

MFPs can require users to be authenticated before accessing MFP functions via the device control panel. MFPs can restrict access to digital sending functions and restrict digital sending e-mail destinations based on the user. MFPs can control access to installed functions and installed applications (e.g. HP Autostore) based on the user. Device usage may also be tracked with associated users.

Integrating MFP access controls with existing enterprise access controls reduces complexity and minimizes administration requirements. HP and its partners support a wide variety of authentication mechanisms, including Microsoft Windows Domain accounts, proximity cards, and smart cards.

HP's Digital Sending Software (DSS) enables Windows and Netware authentication using an intermediary server, while Capella Technologies' VeriUser provides Windows authentication embedded in the MFP. Jetmobile's SecureJet, Ringdale's FollowMe, and SafeCom external authentication each provide smart card, swipe card, and proximity card capabilities.

8.2.2. Network printing authentication

Printers and MFPs may enforce access controls for network printing to restrict usage of devices and the use of high-value consumables. Auditing systems may also use the access controls to log user activity, such as dates and times of documents printed.

The HP Output Server and the Microsoft Windows Print Spooler provide direct integration of domain accounts with printing access controls, which allows control of individual users and groups, including access rights to network printers.

8.2.3. Physical Document Access Control

Documents in the output bin of a network printer are at risk for unauthorized access. PIN and Pull Printing allow print jobs to be saved electronically in the device, or on an external server, until the authorized user is ready to print them. The user provides a simple PIN code, or uses an authentication method supported for other MFP walk-up operations, to release the print job. HP printers and MFPs provide native support for PIN printing, while Jetmobile, Capella Technologies, Ringdale, and SafeCom each provide solutions integral to their authentication products.

Server-based access control: All HP MFPs and digital senders offer server-based Windows NTLM, LDAP, Kerberos, and Novell authentication and authorization that integrates with your existing infrastructure to help your organization manage user access, prevent unwanted printing and digital sending, and help secure access to the management utility to prevent unwanted device configurations. With the exception of the HP 9085mfp and HP Color 9850mfp, all HP MFPs have device-based LDAP authentication (embedded from HP or installable from Capella Technologies). In addition, most HP MFPs have device-based Kerberos available. The HP Officejet 9130 All-in-One supports authentication, as well, via the optional C8267A Secure Digital Sending Solution DIMM. A wide variety of numeric keypad, proximity, and swipe-card solutions are also available, providing a very rich set of capabilities to meet your particular needs.

Color access control: HP's suite of color access control features, available on some HP LaserJet MFPs and printers, lets you closely monitor color use, enable or disable color by individual users or groups or even applications, disable color printing and copying entirely until it's needed for special projects, and report costs back to specific clients, projects, workgroups, or departments.

Control panel lock: This feature within HP Web Jetadmin allows network administrators to deter unauthorized users from changing certain device configurations and control-panel settings by establishing a password and locking the control panel. You can choose from multiple levels of security, locking out specific control panel menus and allowing users to change the rest of the menus, or locking out all of the menus. It is even possible to lock the STOP button.

Private PIN printing: HP MFPs allow a personal identification number to be associated with the print job, which will only be released after that PIN has been entered at the MFP's control panel. Enhanced capabilities, such as retrieval of print jobs at any HP MFP or printer and the use of proximity and swipe cards, can be applied using Ringdale's FollowMe or Capella Technologies' pull printing solutions.

Remote printing security: Secure Document Express provides advanced document-encryption/decryption technology for HP devices equipped with embedded virtual machines. This third-party solution by Capella Technologies provides a fast and economical alternative to certified mail, courier services, and other secure document-delivery methods by allowing users to safely print to any SD-Express-equipped MFP or printer from anywhere on the Web.

8.2.4. HP Secure Erase for Imaging and Printing

To meet the needs for higher levels of print and imaging security, Hewlett-Packard created HP Secure Erase technology for Imaging and Printing. This capability allows the administrator to select how data is erased from storage devices, including print, scan, fax, and copy jobs. Several levels of erase security are provided. The capability is provided as a standard feature on supported HP multifunction peripherals (MFPs), digital copiers, and printers when used with HP's Web Jetadmin (available separately).

HP Secure Erase technology provides a choice of three different modes of erase security, each of which can be configured by an administrator and may be protected from unauthorized changes with a password. The three erase security modes are:

1. Secure Sanitizing Erase mode: Conforms to the U.S. Department of Defense 5220-22.M specification for deleting magnetically stored data. Secure Sanitizing Erase uses multiple data overwrites to eliminate trace magnetic data and also prevents subsequent analysis of the hard disk drive's physical platters for the retrieval of data.

2. Secure Fast Erase mode: This mode completes the erasure faster than Secure Sanitize mode. Secure Fast Erase mode overwrites the existing data once, and prevents software-based "undelete" operations on the data.

3. Non-secure Fast Erase mode: The quickest of the three erasing modes. Non-secure Fast Erase mode marks the print job data as deleted, and allows the MFP's operating system to reclaim and subsequently overwrite the data when needed.

HP Secure Erase technology is applied in two different ways to remove data from storage devices. Secure File Erase erases files on a continuous basis as soon as they are no longer needed to perform the requested function. Secure Storage Erase removes all non-essential data from storage devices in a manner consistent with preparation for decommissioning or redeployment. This operation can be initiated on demand or scheduled for a later date and time.

All data removed from the system by a delete operation is erased using the active erase mode (Secure Sanitizing Erase, Secure Fast Erase, or Non-secure Fast Erase) - this includes temporary files created during the print, scan, fax, and copying processes. User-initiated delete operations, including Stored Jobs and Proof and Hold Jobs deleted through the "Retrieve Job" menu, are also removed using the active Secure Erase mode.

In contrast, the Secure Storage Erase operation will erase stored files even though they have not been retrieved. The HP Secure Erase features will not impact data stored on:

- Flash-based non-volatile RAM that is used to store default printer settings, page counts, etc.
- A system RAM disk (if utilized)
- The flash-based system boot RAM

HP's Secure Sanitizing Erase mode meets the U.S. Department of Defense 5220-22.M overwrite algorithms for overwriting disk files. Using a succession of multiple data overwrites, including the validation of the success of those overwrites, Secure Sanitizing Erase mode can prevent the subsequent physical analysis of the hard disk drive's media for recovery of data. Each byte of file data is overwritten with:

- The fixed character pattern (binary 01001000)
- The compliment of the fixed character pattern (binary 10110111)
- A random character: a stream of random characters is generated using the device's uptime as s seed and is used to overwrite data

To ensure successful completion of the write operation, each overwritten byte is verified.

8.3. Protect Information on the Network

Protecting information on the network insures that network communications between users, administrators, the imaging and printing device, and the workflow are confidential and prevent unauthorized modification by maintaining their integrity.

8.3.1. Network Connectivity with HP Jetdirect Devices

Network connectivity for HP imaging and printing devices is provided by the HP Jetdirect family of products, including internal cards, external boxes, and embedded networking. HP Jetdirect provides many secure network protocols and services, as listed in Table 4-6.

Table 4-6
HP Jetdirect secure network protocols and services

802.1x for Wired Networks	Provides access control to an ethernet network. Network devices that are unable to authenticate to the 802.1x authorization server have all network access denied. 802.1x can prevent unauthorized users from attaching devices to the network as well as insure that only IT deployed and trusted devices, such as those with virus protection software, are allowed access.
IPsec	Allows for strong authentication, confidentiality, and integrity of communications, and can secure network printing and scanning protocols. The HP Jetdirect 635n IPv6/IPsec and Gigabit Ethernet internal print server uses a cryptographic accelerator to provide click-to-clunk performance that rivals unsecured protocols, and supports the IPsec implementations available in all current major operating systems, including Windows, Unix, and Linux.
SNMPv3 and HTTPs	Provide secure management of the imaging and printing device. SNMPv3 provides strong authentication and encryption of management communications and is used by HP Web Jetadmin to provide fleet management of HP imaging and printing devices. HTTPs using SSL/TLS provides security of web protocols and is used for secure management using the device's embedded web server, as well as security of web services such as consumable re-ordering.
Secure IPP (IPP-S)	The secure form of the IPP protocol using SSL/TLS - Secure IPP - requires no additional configuration and is primarily intended for small networks lacking sophisticated IT administration. While Secure IPP may be used in large enterprise environments, IPsec is the recommended protocol for securing printing and scanning functions.

8.3.2. HP Digital Sending Software (DSS)

HP Digital Sending Software allows MFPs to digitally send documents to a variety of destinations, including e-mail, fax, and network folders.

DSS allows the MFP to authenticate a user prior to allowing access to MFP functions. DSS allows integration of authentication functions with Microsoft Windows (using NTLM or Kerberos) and Novell Netware (using Bindery or NDS) operating systems. If authentication is enabled, users are prompted for their username, password, and domain/tree by the MFP. The MFP then transmits these credentials to the DSS server, and the DSS server authenticates the user to the Windows or Novell system as appropriate.

If a remote network folder requires authentication for access, the user's previously provided credentials are used. If the user has not previously provided their user credentials, they are prompted to enter them to access the network folder.

HP Digital Sending Software 4.0 can encrypt scanned documents between the MFP and the DSS Server. The DSS Server may then use the secondary e-mail function to store the encrypted document in a location accessible to third-party applications, such as Omtool, that then securely re-transmit the document to its final destination via e-mail. In addition to the secondary e-mail function, secure sending to e-mail, fax, and network folders may be achieved by securing the network communications between the DSS Server and the remote server using IPsec.

To control e-mail distribution, the SMTP server used by the DSS Server may be configured to enforce internal security policies. Such policies may prevent digital sending to e-mail addresses outside of the internal network or analyzing the content of digitally sent documents to prevent breaches of confidentiality.

8.3.3. Fax/LAN bridging

The analog fax port of an HP imaging and printing device is isolated from the digital network connectivity of the device. Communications to the analog fax are routed directly to the device formatter and cannot be bridged to the digital network, preventing the threat of an attacker connecting to the analog fax through a telephone line and then gaining access to an internal network.

HP is currently in the process of receiving Common Criteria Certification to validate this behavior in the HP LaserJet 4345mfp and 4730mfp.

8.4. Effectively Monitor and Manage

Effectively Monitor and Manage allows for imaging and printing infrastructure maintenance and enables auditing to facilitate compliance with policy and regulatory requirements. Effectively managing network resources is critical to maintaining a secure network.

8.4.1. HP Web Jetadmin for Fleet Management

HP Web Jetadmin (WJA) is the backbone for the administration and maintenance of imaging and printing products, from both HP and its competitors, deployed on enterprise networks. Fleet or batch management enables consistent management and security policy enforcement across a large number of imaging and printing devices. WJA can manage any device that supports the SNMP Printer MIB and allow manufacturers to develop device-specific extensions using plugins.

WJA uses SNMPv3 to ensure authenticated and confidential management of networked devices. WJA allows devices to be manually administered and can automatically discover and configure newly installed devices.

8.4.2. Device and Service Control

Imaging and printing devices support many network protocols and services. Protocols and services that are unused often go ignored, resulting in unintended vulnerabilities, such as unsecured management interfaces or printing protocols that circumvent job accounting controls. HP imaging and printing devices allow individual control over these protocols and services and let administrators enable only the functionality required.

8.4.3. Firmware Updates

Firmware updates can correct product defects and enhance product functionality, and they are an important means for preventing the exploitation of security vulnerabilities. It is important for IT and security administrators to monitor the availability of firmware updates and apply them as necessary. HP releases firmware updates based on the severity of the defect and provides administrators the ability to receive automatic e-mail notifications of releases. HP Web Jetadmin allows an administrator to discover devices using out-of-date firmware and update those devices automatically over the network.

8.4.4. Logging Device Activity

Logging device activities ensures compliance to security and access policies. HP DSS, Capella, SafeCom, and Ringdale each allow device activity, including user, document, and destination, to be monitored. Logging functions can also include configuration and management actions.

Figure 4-12
HP Secure Print Advantage diagram

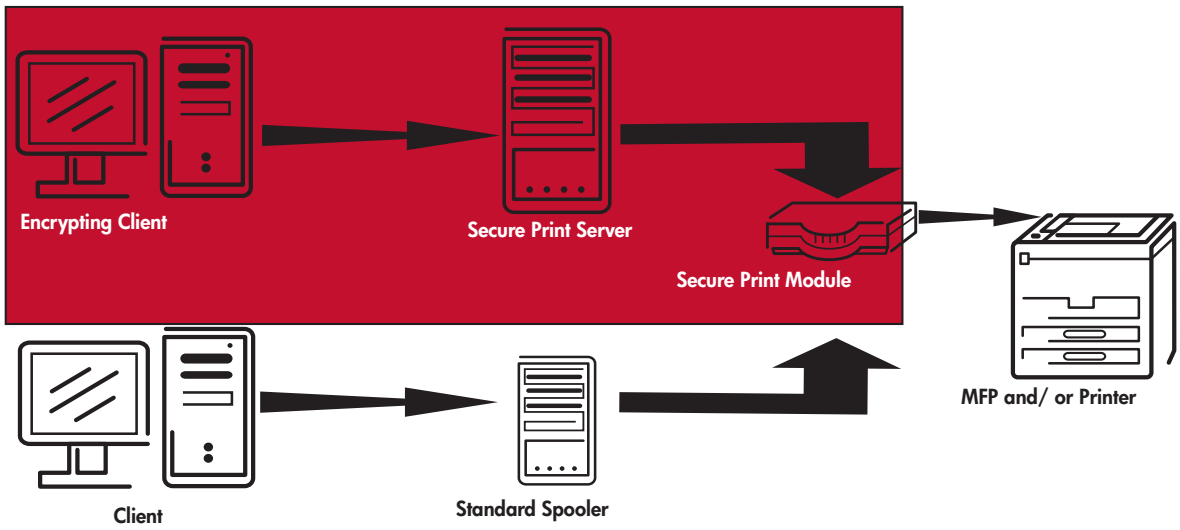

8.5. HP Secure Print Advantage

The HP Secure Print Advantage (SPA) is a comprehensive, end-to-end, architected solution for securing the transmission and printing of sensitive documents and images, without disrupting your existing printer network. It consists of client software, a Secure Document Server, and a Secure Print Module for each networked printer (see also Figure 4-12). It is a traditional print server with policy configuration and enforcement, separation of roles, secure document management, authentication, authenticated audit, security protocol translation, and government certifications.

8.5.1. A Traditional Print Server with Strong Security

HP SPA enables enterprise-wide image and print management for most image and print output devices. It services a variety of clients, including print, fax, and web, and it works with enterprise jobs and queues. Benefits include simplified administration and inventory maintenance.

HP SPA simplifies print network administration in a variety of ways. The single-system interface outputs management resources (print, fax, or web), and a common interface lets you configure dissimilar destination types. The system tracks job and output destination status via graphical display, pager, or e-mail for local and remote administration. By maintaining an inventory of output resources, such as ink and paper, HP SPA also helps you proactively manage supplies.

To support usage optimization and load balancing, the system tracks the source of print jobs and the number of copies.

The HP SPA solution includes three main components: the Secure Document Server, the Secure Print Module, and the secure client application.

- **Secure Document Server**: The Secure Document Server transforms a traditional print server into a certified security appliance. It manages secure printing, performs cryptographic key management, and enforces your organization's security policies. The HP SPA solution meets security standards such as Common Criteria and FIPS 140-2 level 4, the highest level of U.S. government security certification for commercial products.
- **Secure Print Module**: The Secure Print Module decrypts print jobs and manages secure download of updates to the printer. It also secures printer communication management, including user access and authentication at the printer through techniques such as biometrics, passwords, and smart cards. You can install the Secure Print Module with HP LaserJet printers and HP Multi-Function peripherals.
- **Encrypting Client**: The Encrypting Client is a Windows-based client that encrypts data at the client's computer.

HP SPA can overlay an existing network - eliminating the need to reconfigure the network to add security. You can also integrate HP SPA in steps according to a rollout schedule to minimize disruptions.

Simple to deploy and easy to maintain, the HP SPA:

- Introduces security without forcing major network changes
- Works with new or previously installed HP or other printers
- Manages non-SPA print jobs

HP SPA includes policy configuration and enforcement. You can set up policies locally or remotely using dual administrator or security officer control. Rules are enforced according to predefined policies. Device location awareness aids in identifying which type of job can go to which printer (separating secured vs. unsecured areas). And authentication, through PIN and smart card access built into the Secure Print Module, provides additional security.

8.5.2. Authentication, Authorization, and Auditing (AAA)

AAA is important for effective network management and security. HP SPA provides configurable privacy, authentication, and authorization in several interactions, including:

- Client to image and print server communication
- Image and print server to device communication
- User to device communication via smart card or PIN

HP SPA offers multi-layer authentication capabilities including LDAP, Windows Active Directory, and PIN-based authentication.

Auditing capabilities within HP SPA include information about:

- Job sender
- Job completion date and time
- Authentication
- Job deletion, including who deleted it
- User pick up or entry
- Output errors, including partial images or print errors

The HP Secure Print Advantage (SPA) is a comprehensive, end-to-end architecture for the secure transmission and printing of sensitive documents and images. It does not disrupt your existing printer network, but instead preserves assets and allows your organization to work during implementation. With features like policy configuration and enforcement, separation of roles, secure document management, AAA, security protocol translation, and government certifications, HP SPA provides a comprehensive solution for today's print security environment.

8.6. Imaging and Printing-related Certification and Standardization

8.6.1. Common Criteria Certification

HP is currently in process of receiving Common Criteria Certification for Disk Erase and analog fax capabilities for the HP LaserJet 4345mfp and 4730mfp.

HP supports the IEEE p2600's development of an imaging and printing security standard that will allow credible industry-wide Common Criteria Certification and expects to certify products to the standard when available.

While Common Criteria Certification provides a valuable means for assessing the security capabilities of a product, it is important to understand the true significance of Certification, what Common Criteria is and is not, and the role Common Criteria Certification plays in imaging and printing manufacturers' marketing differentiation claims.

Common Criteria Certification provides no credible means for assessing the true security capabilities of hardcopy products today, and should not be used as a measure for purchasing requirements. Common Criteria does not dictate necessary security functionality, it merely provides a means to assess the correctness of a manufacturer's implementation claims.

The varying levels of EAL (Evaluation Assurance Level) certification foster further confusion. Higher certification levels are assumed to provide greater levels of security. However, as certification reflects only the manufacturer's functional claims, the higher levels of certification are frequently meaningless.

The majority of the hardcopy industry currently certifies Disk Erase and Analog Fax functions, but this certification does not accurately portray a product's security capabilities or vulnerabilities. A product may advertise certification of these capabilities while providing no, or rudimentary, protection for the remaining system.

To ensure Common Criteria Certification provides value, it is important to understand the product's complete range of capabilities versus those for which certification is claimed. While certification can prove what a product does properly, it says nothing of what a product does not do, and to what degree that omission represents a security risk.

8.6.2. IEEE p2600

The IEEE p2600 working group is defining a security standard for hardcopy devices, as well as recommendations for the security capabilities of devices when deployed in various environments, including enterprise, high-security, small office/home office, and public spaces. The p2600 working group has broad industry participation, including HP, Lexmark, Canon, Xerox, Sharp, Ricoh, IBM, Epson, Okidata, Equitrac, and Oce. The p2600 standard will provide a means for credibly measuring the security capabilities of individual manufacturers. HP is actively participating within the working group and HP devices support the majority of capabilities specified in the draft documents.

8.6.3. NIST Security checklists

The National Institute of Standards and Technologies (NIST) has been tasked by U.S. legislation to develop checklists that facilitate security configuration of devices likely to be used by the U.S. Federal Government. NIST has requested IT equipment manufacturers to develop these security checklists for their products. NIST will review manufacturer's checklists for relevance and correctness and publish those checklists on a searchable NIST website. Details of the checklist program are available at **http://csrc.nist.gov/checklists**.

HP considers security checklists as a means to significantly improve the security capabilities' ease of configuration for imaging and printing products. A security checklist for the HP LaserJet 4345mfp is available for public review at **http://checklists.nist.gov/repository**, and is currently the only available hardcopy product checklist available from any manufacturer. HP plans to develop additional checklists for hardcopy devices in the future.

8.6.4. Conclusion: Look Beyond Common Criteria Certification

Ultimately, individuals must look carefully at their requirements and not be swayed by manufacturer advertising claims. Common Criteria Certification adds significant cost and development time to products, while providing limited assurance to the product's actual capabilities and potential vulnerabilities. Products that are not certified may actually provide more robust security capabilities than products that are certified. NIST security checklists simplify the complex process of enabling security functions, and better illustrate the product's capabilities.

8.7. Conclusion

HP imaging and printing has evolved with enterprise security needs. HP offers imaging and printing devices with a broad range of security capabilities, including high-security products that allow operations in the most demanding environments and the tools to effectively manage large-scale deployments of those devices.

While it would be impossible to prescribe all of the security requirements for an enterprise's imaging and printing environment, the following recommendations may be used as a starting point for enabling that security.

1. Assess Common Criteria Certification needs: The features being certified by the hardcopy industry are not representative of the true risks that face imaging and printing devices. It is critical to scrutinize certification and assess the capabilities of the device against actual needs.

2. Fleet/batch manage using HP Web Jetadmin: HP Web Jetadmin provides consistent management of enterprise-deployed imaging and printing devices and is critical for maintaining a secure environment. Fleet management aids in the consistency of policy enforcement and assists in audit and regulatory compliance.

3. Update firmware images: Firmware updates protect against product defects and vulnerabilities. HP provides automated firmware update notification services, and HP Web Jetadmin aids in deploying updates across enterprise environments.

4. Disable unused ports and services: Frequently, imaging and printing devices have unused capabilities that are enabled. In some cases, these capabilities may enable functionality counter to the intent of the administrator, such as leaving insecure management protocols accessible, when only encrypted management is desired.

5. Implement access controls: HP printers and MFPs allow a variety of user-level authentication mechanisms, including passwords, proximity cards, and smart cards. Access controls can ensure that only authorized users utilize the imaging and printing infrastructure, while authentication capabilities provide assurances of who is using the environment, and how they are using it, which aids in audit and regulatory compliance.

6. Implement secure protocols: The sophistication necessary to sniff network traffic has been reduced by the distribution of hacking tools, as well as by legitimate network analyzers. IPsec secures existing printing and scanning applications with strong encryption, while SNMPv3 and HTTPS secures management functions.

9. HP Trusted Infrastructure Services

HP offers a wide range of services capability for designing and implementing trusted infrastructures that meet organization's business needs. HP's Consulting and Integration Services combined with HP Technology Services and HP Outsourcing Services offer trusted infrastructure services at every point in the security lifecycle. The following is an overview of HP's Trusted Infrastructure services (see **www.hp.com/go/security**):

- Infrastructure review and implementation design
- Security assessments across the infrastructure
- Physical asset protection
- Network, system, and host security
- Adaptive Network Architecture
- Application security and application auditing
- Security workshops and training

10. Trusted Infrastructure Summary

As reliance on IT infrastructures increases for businesses and society, we face important challenges. We must stay ahead of the security needs for reliable infrastructure technologies. Fundamental IT building blocks must be innovated and redesigned to include security features. From clients to servers, from networking to storage, and in printing technologies, infrastructure security mechanisms must be continually improved to support adaptive and flexible IT solutions.

HP is investing to ensure that we continue to deploy secure and reliable trusted infrastructures. HP is an industry leader, driving this agenda across platforms, OSs, and infrastructure solutions. Importantly, HP's leadership in the TCG has brought the industry together to greatly increase baseline security of infrastructure technologies to meet current and future customer needs.

Alongside other efforts, such as establishing secure development practices within HP and driving infrastructure technology standards, Trusted Computing provides the security building blocks that allow the IT industry to continue to innovate and deliver the power of IT across reliable trusted infrastructures.

Table 4-8
HP trusted infrastructure solution offering summary

Trusted Infrastructure Component	Solution	URL
Network Security	HP ProCurve Networking HP Virus Throttle HP Adaptive Network Architecture HP Proliant Essentials Intelligent Networking Pack HP IPFilter/9000 HP ProLiant DL320 Firewall/VPN/Cache Server	www.hp.com/go/security/trusted Click solution components tab
Host Security	HP ProtectTools HP-UX 11i Linux HP NetTop HP OpenVMS Tru64 HP NonStop Systems HP Atalla Security Products HP Trusted Compliance Solution for Energy HP Application Security Center HP Enterprise Mobility Suite	www.hp.com/go/security/trusted Click solution components tab www.spidynamics.com www.hp.com/go/ems
Storage Security	HP StorageWorks LTO-4 Ultrium 1840 Tape Drive HP StorageWorks Secure Key Manager HP Compliance Log Warehouse	www.hp.com/go/security/trusted Click solution components tab
Imaging and Printing Security	HP Secure Print Advantage	www.hp.com/go/security/trusted Click solution components tab
HP Security Services	Trusted Infrastructure Services	www.hp.com/go/security/trusted Click solution components tab

For additional information, refer to the following resources:

Trusted Computing Platforms: TCPA Technology In Context, by Dr. Siani Pearson et al., Prentice Hall PTR, July 2002, ISBN 0-13-009220-7 (Order at **www.hp.com/hpbooks**.)

Cisco's documentation web page: **www.cisco.com/univercd/home/home.htm** (For general references, see the links under the "Hot Items" and "Networking Information" headings.)

SANS Institute Information Security Reading Room: **www.sans.org/rr/** (An excellent selection of white papers on a wide variety of network and general security topics)

Inside Network Perimeter Security, Stephen Northcutt et al., Sams, 2nd Edition 2005, ISBN 0-6723-2737-6

Information Security Management Handbook, Harold F. Tipton and Micki Krause, Auerbach Publications, 5th Edition 2004, ISBN 0-8493-1997-8

Official (ISC)2 Guide to the CISSP Exam, Susan Hansche, John Berti, and Chris Hare, Auerbach Publications 2004, ISBN 0-8493-1707-X





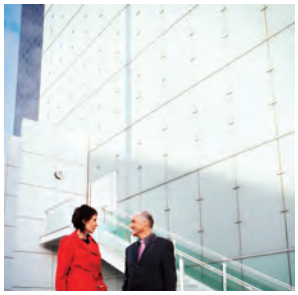






Chapter 5
Innovation in Information Security

"HP Labs' corporate immune system technologies transformed debilitating attacks, which have caused widespread interruptions within other companies, into localized annoyances at HP. This secures our Adaptive Enterprise. It is HP Invent at its best."
-Sherry Ryan, HP Director of Information Security

Introduction

The introduction to this handbook has outlined the changing nature of enterprise computing, the rapidly changing threat landscape, and the consequences for IT security. Throughout the handbook, we have described HP's current offerings that help our customers achieve appropriate IT security. In this section, we focus on the longer-term future and the contributions of HP Labs, HP's central research organization. HP Labs' function is to deliver breakthrough technologies and technology advancements that provide a competitive advantage for HP, and to create business opportunities that go beyond HP's current strategies.

In this future, businesses will rely on highly mobile, flexible, shared infrastructures, supporting rapidly changing business processes. In addition, personal experience-driven trends such as Web 2.0, which empower users to share information more richly, will shape the way employees manage information. This will place new challenges on how enterprises manage information. There will be many challenges in keeping businesses and business processes safe and secure. However, alongside these challenges, we also see opportunities to create qualitative changes in IT security that will be key enablers in making this new world a reality.

As context it is useful to think about the lifecycle of security management. As shown in Figure 5-1, security starts with an analysis of the IT-associated risks. From such analysis, policies and controls evolve that shape the way security investments, operations, and configurations are executed. This in turn leads to mechanisms that enforce policy in the infrastructure, and finally to the monitoring of key parts of the overall system, including incidents and events, that get aggregated and analyzed to provide assurance that regulations are being met and risks are appropriately mitigated. This all takes place in the context of a changing threat landscape, where we expect to have to protect against more targeted attacks and the exploitation of increasing numbers of unknown vulnerabilities.

Figure 5-1
Security management lifecycle

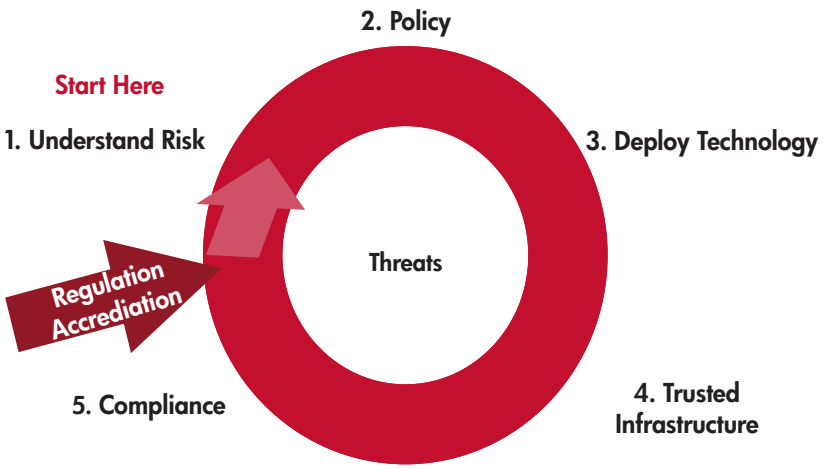

Many enterprises are consolidating or outsourcing their data centers. At HP Labs we are investing beyond this stage to when data centers are shared and federated, creating a true compute utility. This program involves virtualization, automation, and even integrating smart cooling. Such an agile infrastructure will enable businesses to change IT more rapidly, but each change will have risk implications, which in turn means we have to get more efficient and effective at operating this security lifecycle. Today's best practices and mechanisms will all help, but at some point soon a qualitative change in capability will be required.

This is a tremendous challenge, and one that we must overcome to realize the future IT vision. Today people and manual processes control much of this lifecycle, using crude tools and infrastructure that lacks the necessary security properties. HP Labs is creating technology to model, connect, simplify, and where possible, automate the various activities. We are also investing heavily to create trusted infrastructure with better security properties that will ensure that policies, remotely passed to shared and virtualized components, will be reliably enforced and meaningful, trustworthy assurance information will be returned.

This chapter provides a short overview of some of these ideas, covering the Economics of Information Security, Identity Management, Trusted Infrastructure, Assurance and Threat Management.

Unique among technology companies, HP has a broad and deep set of offerings across all market segments: from consumer to enterprise, from small and midsize businesses to the public sector. Reflecting this diversity, HP Labs' research portfolio includes projects in imaging and printing, advanced architectures, mobile systems, nanotechnology, business intelligence, and media systems. Many of these technologies have the potential to disrupt trust and security. Three examples that we picked out to cover in this chapter are Quantum Cryptography, Memory Spot and Trusted Printing.

1. Trust Economics

Creating a security architecture, particularly one which meets the security governance objectives of the business within the overall requirements of IT governance, is a challenging task. While senior managers or officers of an organization use quantitative measures for financial risk management in corporate governance, they typically resort to a qualitative understanding of IT risks to manage information system risk. Developing a quantitative information system risk management toolset is a grand challenge in information security, which HP Labs is tackling through its Trust Economics program.

Trust economics is the conceptual framework that HP is developing to pursue the study of information security policies, protocols, and investment strategies. HP's perspective is one of systems thinking, embracing studies of both the economic factors and user behaviors that bear on IT risk management. Two key problems facing senior managers with responsibility for information and systems security are:

- Developing an economic understanding of how to formulate, resource, measure, and value security policies
- Understanding the attitudes of users to both information and systems security and of their responses to imposed security policies

A model of the system and its economic environment is necessary to assess the effectiveness and value of security investments. A rigorous understanding of the behavior of users, together with the economic value of the system's security measures, can be captured within an extension of some established mathematical systems modeling techniques. Our technique integrates the following three perspectives:

- Modeling the behavior of the users of systems, both internal (operators, staff) and external (customers, regulators), in the context of security policies and protocols
- Mathematical modeling of systems, organizations, and networks, including the security policies and protocols that govern access
- Economic modeling of the costs and value of security policies, protocols, and technologies

The main challenge here is to understand how to develop and integrate effectively two different kinds of modeling. We must extend the mathematical modeling of the technological aspects of a system to encompass the users of the system; we must also integrate economic models as valuation methods. A significant challenge is to build models at levels of abstraction that capture just the questions of interest and avoid irrelevant, complicating details.

Mathematical systems' modeling uses methods drawn from algebra, logic, computation theory, and probability theory. User modeling uses psychological models (e.g., cognitive architectural and knowledge models) based on our understanding of humans. Embedded user models (those held by systems representing characteristics of users) often rely on statistical methods. Assessing the validity of such models requires empirical study, either in the form of field or ethnographic studies, and/or experimentation. We aim to integrate these approaches into a suitable economic model and develop a new science of systems security services.

Example topics we are investigating include:

- Establishing effective security cultures: By modeling the behavioral consequences of policy choices, we seek to establish mechanisms for selecting those choices that promote more effective “cultures”. For example, by providing security consultants and systems engineers with the tools and techniques to quantify decisions, better understood and more justifiably trusted systems can be built.
- Employee risk assessment: By establishing mechanisms for assessing the consequences, relative to a given security posture of particular patterns of behavior, we will be able to provide a framework to assess the possible security implications of particular policy and implementation choices relative to the intended users of the system.
- Investment strategies for information security: Understanding the investment options against the dynamic threat environment.

2. Identity Management

In many ways identity management is already a mature discipline, with many good tools and best practices, as described in the identity management chapter (Chapter 3) of this handbook. However, the problem of identity and access management continues to get more complicated. Solutions need to scale to a growing number of users, roles, data items, and resources. Constant change, fueled by de-perimeterization and increasing numbers of acquisitions, virtual organizations and transient partnerships, demand techniques that simplify identity management.

Two trends create new challenges. The first challenge comes from the changing nature of resources and content that must be protected, and the second is the devolution of identity control from the organization to the individual. Content is getting more complicated, whether it is highly structured in centrally controlled data warehouses, semi-structured formal documents, or even ad hoc items on collaboration portals or client machines. Personal experience-driven trends such as Web 2.0 empower users to share information more richly and place the individual in greater control of information and identity. Businesses must adapt to these changes while preserving the control that allows them to meet regulatory requirements and manage business risk.

This section presents HP research addressing two related challenges, in the context of policy enforcement/control and policy planning/definition:

- Content-aware Access Policies: Access policy expression and enforcement to help deal with complexity, scale and assurance
- Role Discovery: Sophisticated algorithms we have developed to "discover" implicit roles and thereby help plan major Identity Management changes

2.1. Content-aware Access Policies

Enterprise information can come from a number of forms including database tables, semi-structured documents (e.g. XML), or even unstructured files and volumes. Moreover, the source and manner that this content creation occurs can vary from formal business processes, ad hoc collaborations, or individual content generation (e.g. blogs). Finally, the subject of this information can span employee data, financials, sales collateral, contracts, intellectual property, and customer, partnership, and forecasts data, each with subtly different business sensitivity and regulatory control.

The challenge is how to maintain the ability for the business to extract full value and flexibility from this data, while maintaining necessary control. As a simple example of the problem, consider access control rules to data containing personally identifiable information (PII). Privacy protection legislation demands that organizations clearly state the purpose for which they will access and use data containing PII. There may be much legitimate value that a business can make of such data, but without ways to differentiate usage, it is difficult to use mechanisms to allow, anonymize, or restrict access in line with PII regulations.

HP Labs has explored context aware access policies to deal with this. The goal is to move beyond simple binary access decisions. Content-aware access policies take account of the content or context of a request for data and might return limited portions of the data. An early example of this is the HP Labs "privacy policy enforcer", which specifically takes account of purpose before granting access to PII data. A richer example is an engine that takes a structured (e.g. XML) document with policies potentially associated with each component of the document. For each request, the engine uses credentials, context, and policy to create a valid view of the document that it returns to the requestor.

This approach allows fine-grained access policies for dealing with anything from enterprise blogs to patient records. For example, it becomes possible to express and enforce audited access for doctors accessing a patient's records remotely or outside normal hours; also, partial access could be granted to nurses, hospital administrators and even researchers with particular credentials (e.g. to study certain kinds of diseases).

Richer policy is only part of the problem, there are still problems to make it easy to have content with appropriate structure, and as IT environments become more distributed and shared we will need infrastructure with better security properties to ensure that these richly expressed polices are being enforced.

Individuals have a similar problem in their use of information and communication technologies in their private lives, and this is compounded by their using multiple partial identities for different online relationships and purposes. HP Labs is researching the use of a policy-driven approach for managing personal information within personal client devices and controlling its release to other parties based on their policies, reputations and other trust factors.

2.2. Role Discovery

Traditionally, IT personnel manage access rights directly. However, in large enterprise environments, this becomes impractical simply due to the scale and dynamic nature of the problem. In a large organization, it is common to have tens of thousands of users connecting to a similar number of resources. Role-based Access Control (RBAC) is a standard approach whereby an intermediate set of entities, called roles, are used to aggregate resources. For example, a role might be defined for an accounts receivable clerk to grant access to the set of resources that such a person would need in order to do their job. Then users can be simply assigned to roles. This greatly simplifies the management problem and is particularly effective because the approach more closely aligns to the business objectives of the organization.

A major challenge is transforming an organization from a traditional access control system to an RBAC system. This labor-intensive process requires an organization to initiate a role development study in which the roles need to be researched and meticulously defined to meet the organization's business needs.

At HP Labs, we have developed a new approach, called role discovery, to make this role development process more efficient by discovering roles that are inherently defined in the organization's existing traditional access control environment.

The technical innovation behind role discovery is the formalization of this problem in terms of graph theory. We can show that a set of traditional access control rules can be represented as a kind of graph called a bipartite graph. Moreover, the transformation of that system to a set of RBAC rules corresponds to transforming that bipartite graph into another kind of graph called a tripartite graph.

This is a well-known problem in theoretical computer science. It is a particularly difficult problem, for which no known algorithms can guarantee to find the optimal solution in a reasonable amount of time. In fact, it is even difficult to find an approximate solution efficiently. However, we have developed some algorithms that work extremely well, in practice, on real data. These algorithms scale to very large problems and are very fast.

HP's internal IT department is using role discovery to help simplify the way we manage how external business partners connect into internal HP systems. We are defining a new network access control paradigm that leverages the simplicity and manageability of RBAC for the network layer. Role discovery is a tool to help make that transition more efficient.

3. Trusted Infrastructure

Utility computing introduces extraordinary flexibility for enterprise computing. It will mean compute resources available on tap with tightly defined and easily expressed parameters. For example, an enterprise service might be configured using the following specification:

I need resources to run an ERP application. The service needs to be globally available with load/performance requirements varying across time zones. Given that it is running in a shared environment, in addition to the normal security operations, I need extra assurance that confidentiality and integrity of my resources will be maintained.

Virtualization is a key ingredient to realizing this vision. The ability to virtualize computing platforms enables us to run multiple operating systems on a single platform. This is useful as many legacy applications can now share physical resources, consolidating servers, and creating the flexibility to move applications around as demand changes. However, securing a single platform is challenging and something that must exist as a foundation to the security of the management system controlling these virtual resources.

Standard usage of virtualization has a virtualization layer that restricts the view the guest operating system (virtual machine) has of the physical platform. This can be extended to restrict the access each virtual machine has to other infrastructure elements such as the network, and shared storage. This is very useful for utility computing as it enables us to create virtual (and flexible) network and storage infrastructures that run over shared physical network and storage resources. However, this makes the virtualization layer a critical point for security. Any defects here could compromise the security of virtual networks/storage, the management system, as well as the virtual machines "owned" by the customers of the utility.

Today's virtualization platforms typically run a full-blown operating system directly. This is a large and complex piece of code to rely on to isolate functionality and enforce policies. From a security perspective, it is desirable to rely on the much simpler (and smaller) virtual machine monitor (VMM) which is mainly responsible for the lifecycle and scheduling of virtual machines. The problem is how to design a management system for such a virtual platform which itself is not subject to all of the vulnerabilities of a large-scale operating system.

HP Labs is pursuing research to create a trusted virtual platform. One component of this research is that we would like to use the VMM to isolate various security services from the main management system. However, this is not practical or secure if each service runs in a large-footprint OS. As such for the Xen virtualization platform, we have built a very light-weight library OS with communication services, suitable for running a number of security services. This enables us to create more trustworthy security services such as component identity, integrity and audit.

A second problem that quickly emerges with all these virtual components is how to ground or root trust. If a virtual software component presents itself, on what basis can it be trusted? Trusted computing, as defined by the Trusted Computing Group (TCG), provides a physical root of trust for identity and attestation that is a natural answer to this. It is possible to create a chain of trust through the virtualization layer so that identities and attestations presented by virtual components are rooted in the physical TPM. This means that a remote machine (customer or management utility) can verify that each component is running the expected software, and that it is running on a virtual platform that can be trusted to enforce security policies, and the remote machine can isolate the component from other software sharing the same physical platform.

In addition to securing platforms for next generation data centers these security virtualization properties apply to client machines. For example, with these mechanisms, users can trust using one virtual machine for playing games, another for private banking, and another to access their workplace intranet. Moreover, the bank and the employer can use TCG-based attestation to gain assurance that the correct virtualization software is running, and that bank and corporate approved software and configurations operate in the virtual machines with which they are interacting.

HP is leading the EU-funded research collaboration "Open Trusted Computing" (see **www.opentc.net**), which seeks to combine open source, virtualization and TCG mechanisms to create trusted platforms. This project has already demonstrated the banking example described above.

4. Assurance

While it is important to build security enforcement mechanisms into systems, it is equally important to build in the hooks to know the security is working and to detect when users are misusing systems. As we build compute utilities with automated management, we need systems to provide automated assurance. Today because so much security relies on best practices, assurance is largely manual, with auditors using spreadsheets to reconcile events with process controls. There are two challenges here: first, how to build a framework to automate this kind of reconciliation, and second, to ensure there are independent paths in the systems and infrastructure that can be relied on to provide trusted information to this framework.

The work on model-based compliance described in the Governance and Compliance chapter (Chapter 1) addresses the first of these problems. The modeling tool allows auditors and risk officers to create intuitive graphical representations of the controls and how they should be tested. The models also have precise semantics, which enables the structures to be integrated directly with data from the IT environment. Through customer pilot projects, this approach has been shown to remove much of the routine data collection and analysis, and allow more of auditor's time to be devoted to more meaningful risk discussions.

We can apply this approach to current environments, but today designing the data collection is a manual integration project. In utility environments, where the infrastructure is always changing, the models need to adjust automatically to collect data appropriate for the current configuration. Moreover, because multiple parties share the infrastructure, we have to address the second problem of making the collected data trustworthy.

This is where our work on assurance and trusted infrastructure come together. In the trusted infrastructure subsection, we hinted at how the trusted virtual platform design enables us to build small-footprint security services that can be isolated (i.e. independent) from other parts of the management system. This is the basis for the work we are doing to create a secure and independent audit service which can be deployed (perhaps multiple times) onto a physical platform. The components of the audit service, attested to using TCG-based mechanisms, provide configurable and independent information monitoring of the basic building blocks of the utility.

Significant challenges remain, however, in determining what mixture of events, configuration, and state attestations provide adequate assurance about a virtual platform.

5. Threat Management

In the film *The Matrix,* the main character faces a choice: swallow a blue pill and continue to experience "normal" life, or take the red pill and discover new and deeper realities about the world.

In many ways, the majority of security activity - patching, protecting or monitoring against known threats - can be seen as acting in the blue pill world. This world assumes that vulnerabilities are discovered, patches are created, there is a race to deploy patches before exploit code can get to the vulnerable systems, and that this race is getting harder and harder to win. The deeper reality or truth is worse than this: while it is known that attackers actively discover vulnerabilities, the software vendors do not necessarily learn of all the newly discovered vulnerabilities, even after an attacker exploits them.

The challenge is to create defensive techniques that address classes of threats without the need for knowledge, or even the existence, of specific threats at any point in time. The trusted infrastructure research program, which can protect vulnerable components through isolation, is one approach to this. Another more direct and novel approach HP Labs is taking involves exploiting the differences between information and its representation by specific data or data formats. For example, rather than checking data for known patterns or signatures of exploits, we can often change the data representation of a piece of information while preserving the value in the information. Take Figure 5-2. Can you spot the difference?

Figure 5-2
Threat management example Figure 1

How about when looking at Figure 5-3?

Figure 5-3
Threat management example Figure 2

```
ff d8 ff e0 00 10 4ª 46 49 46 00 01 01 01 00 48 |......JFIF.....H|
00 48 00 00 ff fe 00 36 20 49 6d 61 67 65 20 67 |.H.....6 Image g|
65 6e 65 72 61 74 65 64 20 62 79 20 45 53 50 20 |enerated by ESP |
47 68 6f 73 74 73 63 72 69 70 74 20 28 64 65 76 |Ghostscript (dev|
69 63 65 3d 70 6e 6d 72 61 77 29 0a ff db 00 43 |ice=pnmraw)....C|
00 01 01 01 01 01 01 01 01 01 01 01 01 01 01 01 |................|
01 01 01 01 01 01 01 01 01 01 01 01 01 01 01 01 |................|
01 ff db 00 43 01 01 01 01 01 01 01 01 01 01 01 |....C...........|
01 01 01 01 01 01 01 01 01 01 01 01 01 01 01 01 |................|
….
….


89 50 4e 47 0d 0a 1a 0a 00 00 00 0d 49 48 44 52 |.PNG........IHDR|
00 00 00 ab 00 00 01 0d 08 02 00 00 00 ff b8 02 |................|
09 00 00 00 09 70 48 59 73 00 00 00 48 00 00 00 |.....pHYs...H...|
48 00 46 c9 6b 3e 00 00 00 09 76 70 41 67 00 00 |H.F.k>....vpAg..|
00 ab 00 00 01 0d 00 d2 65 48 5e 00 00 80 00 49 |........eH^....I|
44 41 54 78 da e5 bd 79 9c 1b f5 79 3f fe 48 a3 |DATx...y...y?.H.|
19 69 46 d7 4a 2b ed 7d d8 92 c1 7b 78 8d ed 5d |.iF.J+.}...{x..]|
6c 0c 3e 93 35 60 48 4c 08 01 92 86 2b 4d 7f 86 |l.>.5`HL....+M..|
26 69 1a 92 b6 a6 49 db 24 cd 65 7f db 40 d2 1c |&i....I.$.e..@..|
….
….
```

The information conveyed by both images in Figure 5-2 is seemingly that of the same tiger, however looking at Figure 5-3, you can see that the constituent data representation of the images are significantly different: one image is in JPEG picture format and the other is in PNG picture format. Since exploits have to be extremely precise to utilize a vulnerability, it is likely that any exploit will be disabled by such a data conversion, and the conversion does not adversely affect the user. In this example, converting from the JPEG to the PNG picture format doesn't change the picture's visual image but does change the data format. If there was a hidden exploit for JPEG image rendering software, it would be lost in the transformation into the PNG format.

More generally, there are a vast number of equivalent data formats available. We have used these to create a service that blindly and randomly transforms data, while maintaining information equivalence. The service disrupts any "bad" data while not affecting the information provided by "good" data. That is, the user still experiences the value from the file, and any malware or vulnerability has been cleaned out.

Clearly, there are more complications if the user wants to do more than view information, for this we are exploring notions of functional equivalence. The approach may not be applicable in all instances of data (for example, executables present quite a different challenge). However, a powerful aspect of this approach is that it protects against unknown attacks.

6. Quantum Cryptography

Quantum computers are uniquely capable of factoring large numbers and this has the potential to disrupt many assumptions made about current cryptography. This fact has motivated the HP Labs research on Quantum Key Distribution (QKD), which provides a means for two parties to generate secure, shared-secret material, which could be used as one-time secret keys or pads to encrypt and decrypt information.

Used properly, one-time pads provide guaranteed, unbreakable cryptography, even with advances in quantum computation. The rules of quantum physics guarantee that using QKD makes it impossible for an eavesdropper to snoop on an interaction undetected.

The one-time pads generated using QKD may be used in a number of ways to protect e-commerce, or to identify individuals to each other. In addition, one-time pads provide the security required for performing the QKD operation.

Other research and product groups offer QKD systems. The problem that HP Labs is addressing is how to make this technology available to the mass consumer market.

We have built a low-cost, free-space quantum cryptography system using off-the-shelf components that is able to generate and renew shared secrets on demand over a short range (up to one meter) in shaded daylight conditions. The transmitter uses a compact diffraction-grating optical element design, which we plan to incorporate into a hand-held device such as a smart card or mobile phone.

As an example of the problems to be solved, imagine a very weak light signal (small numbers of photons) being received amongst a background of random light. The signal has to be identified and extracted. Once extracted, the endpoints have to communicate in order to error-correct and generate a shared secret. A full software system has been developed to handle this signal processing.

Currently the system can generate around 40,000 bits of secret keys from a one-second interaction between transmitter and receiver, depending on the light conditions.

Figure 5-4
Micrograph of the 32k memory spot chip

7. Memory Spot Technology

HP Laboratories has developed a miniature wireless data chip called memory spot, which at present has no equal in terms of its combination of size, memory capacity and data access speed. Embeddable into physical media (e.g. paper), this technology has the potential to completely change many security rules and assumptions. For example, this technology can be applied to allow an organization to securely link digital and physical objects such as a file and a printout of that file.

Memory spot is a fully functional chip that, in its current design, is fabricated onto a 1.4mmx1.4mm square of silicon. It is a near-contact technology that allows a very fast data transfer rate of 10Mb per second. This is 25 times faster than the current, fastest RFID system - and this higher transfer speed permits rapid download of large amounts of data in a very short time; a "touch-and-go" style of interaction is possible even for large data transfers. For example, a 20-second audio file can be transferred from a memory spot device in under 50 milliseconds.

From a security angle, each memory spot device is equipped with an on-chip challenge-response authenticator based on an industry standard SHA-1 algorithm. This feature is useful when local authentication is required or in situations where local authentication is the only available option. If deployed correctly, memory spot's on-board authenticator makes the chip virtually impossible to clone.

Each memory spot device is also equipped with a rudimentary yet effective on-board data read access mechanism based on a 224-bit password. If activated, data can be added to a memory spot but cannot be read unless supplied with the correct password. This feature enables a memory spot device to act as a data carrier without the need to review any shared secret or public key.

Memory spot has application in areas such as pharmaceutical anti-counterfeiting, assuring the provenance of high-value items (e.g. aviation parts), and validating documents such as birth and marriage certificates.

8. Trusted Printing

HP has unique and significant experience in imaging and printing. We complete this security innovation chapter by drawing out two final examples where HP Labs have made contributions to create trusted printing solutions.

Counterfeiting of products and documents is an illicit industry costing legitimate brand owners hundreds of billions of dollars each year. Two new businesses are forming within HP: one providing product security and the other providing document security.

The first business is called Trusted Track and Trace. Product Track and Trace is the generation and storage of a provenance record for a given packaged good. This includes the package identifier along with information about where and when it was scanned while in transit from the manufacturing site to the end user. Bar codes or RFID, increasingly using the Electronic Product Code (EPC) standards with EPCglobal certification, are the standard mechanisms for Product Track and Trace. Trusted track and trace augments the EPCglobal traceability through a security label that incorporates multiple features that are both difficult to copy and difficult to reproduce without using the HP Indigo variable data press. Variable microtext and wide-gamut color features are combined with EPCglobal-compliant bar coding to create an eye-catching, authenticable security label.

This program has benefited from the HP security community's expertise in consulting and integration of EPCglobal-compliant solutions. The HP security community has also provided leadership on security issues for authorization, data retention, database access control, and other areas. This program has also benefited from HP's security application threat analysis methodology: a full threat analysis review was performed for the Trusted Track and Trace program during the design phase to ensure that system-level security threats were addressed in the final approved product design.

The second new business is called Trusted Hardcopy. This business solves the customer need of being able to trust documents that enable high-value transactions. Financial, government and educational institutions, among other document-issuing authorities, create documents that are later used by banks, real estate offices, insurance offices, etc., to approve transactions at points of authentication. These transactions of value are subject to a high rate of document counterfeiting. HP's Trusted Hardcopy solution provides document protection at document creation, on the document itself, and at the site of authentication. The document contains security deterrents, including a 2D barcode and copy/tamper-evident features, which combined provide a hash of the salient (indexing) information on the document. These features are generated (hashed) and affixed to the document during its creation, and then can be read and validated at the point of authentication.

This program has also benefited from the HP security community's strength in secure printing. Secure printing provides authorized printing, where a print job is not completed until authorized by the requestor (e.g. with a password, smart card, etc.) and where the document printed is securely removed from the print device's memory afterward.

9. Conclusion

We expect the context to remain challenging; businesses will continue to balance the demands to reduce cost and maintain control of IT. In addition, they will quite rightly seek new ways to improve business processes and gain competitive advantage through changes in IT.

This short chapter has described just a sample of the security research in HP Labs. Each of the subsections has shown, in different ways, how we can use science, mathematics and technology to help businesses operate the security lifecycle more efficiently or with more agility. More details can be found in the HP Labs technical reports series at **www.hpl.hp.com/techreports**.

The context is always changing which causes the research to evolve continually. Therefore, we suggest if you want an up-to-date view on our research, you explore the HP Labs web pages at **www.hpl.hp.com/research**.

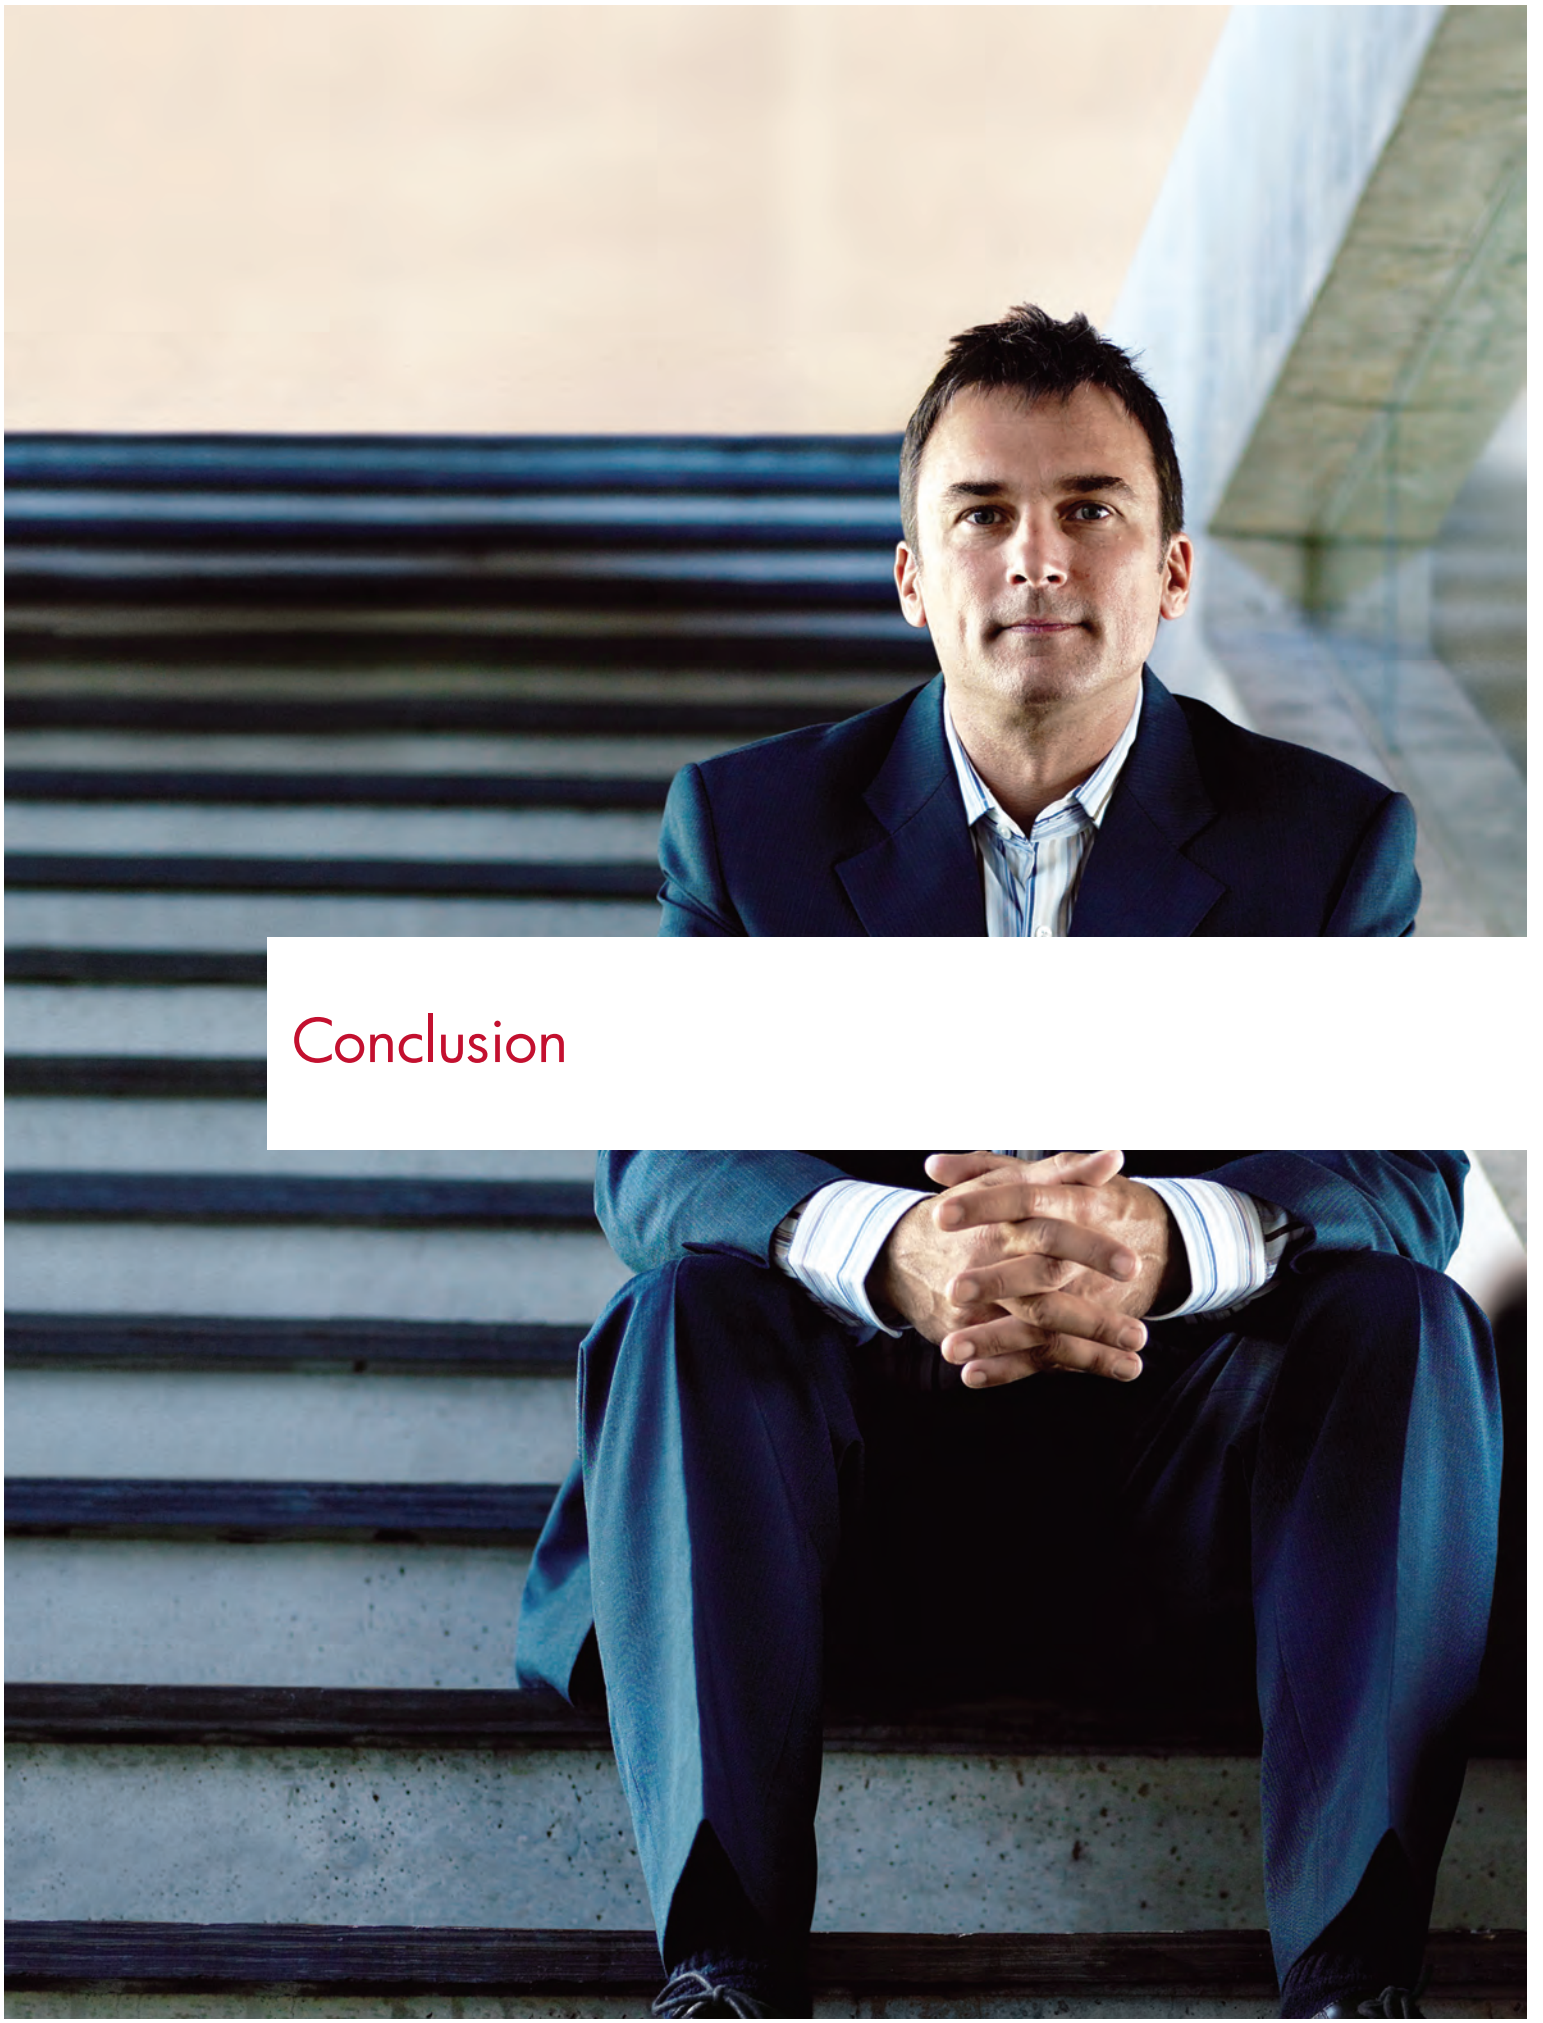

Conclusion

"HP's strategy is to build security into its products, drive industry standards on security and privacy, align business and regulatory requirements with security lifecycle delivery, and innovate ways to deliver a safer IT environment for our customers."
-Tony Redmond, Vice President, HP Security Office

Inventive, Reliable Security

All of HP's businesses sell products, services, or solutions that require varying levels of security, to be both acceptable to customers and competitive in the market. Security is increasingly becoming an attribute that is associated with quality. HP wants our products, services, and solutions to be secure in operation, deployment, and use. We want to be known as a company that designs for security and privacy, drives best practices, contributes in a significant manner to new security standards, and delivers a safer IT environment to our customers. We draw on our own resources as well as those of our pure-play security partners, such as Symantec, Check Point, Nokia, Cisco, and VeriSign, to deliver hardware and software products and services that contribute to our strategy and meet the requirements of our customers.

Figure 1
HP's security framework

Delivering a safer IT environment requires a framework for rapid and effective response to threats and business objectives. HP focuses on the key areas of governance and compliance, identity management, proactive security management, and trusted infrastructure to bring our customers the safe, proactive, and adaptable IT environment that is necessary to support the objectives of companies and organizations today and in the future.

Governance and Compliance

One of the most striking features of today's business environment is its dynamic nature. Successful companies capitalize on change, turning what is often unexpected and disruptive into a business advantage. HP's Security Governance Services provide companies and organizations with an enterprise-wide policy foundation, a governance model, and an organizational structure. The program can apply to the entire enterprise or to a business line, and it defines the integration and orchestration principles that shape the enterprise security system. This program meets the requirements of the ecosystem in which the enterprise operates, including regulations, business community practices, technology constraints, and the culture specific to the enterprise. Security governance provides guidance on how IT staff translates business security requirements into security measures and implementations. HP's Security Governance Services include a broad set of offerings delivered across the governance lifecycle to build an enterprise-wide policy foundation, a secure and agile architecture, process framework, and an organizational structure. Together these services enable dynamic businesses to manage the risks associated with their information assets.

Proactive Security Management

It is often joked that the most secure computer is one that is in a guarded, locked room…and is also turned off. The point of the joke is that there is no such thing as 100% security and the most secure system is one that is not useful. The reality is that there is a set of trade-offs or variables to manage - such as costs, asset values, security technologies, people. Proactive security management is the science of managing those variables - with people, processes and technology - to support an organization's goals, and do so while maintaining an acceptable level of risk. The environment for our IT infrastructures includes an ever-changing state of threats, an evolving set of vulnerabilities and the basic, human-nature condition that if something has value then there is at least one person who might try to take it.

Security management has matured far beyond simply keeping out intruders or presenting a single console to coordinate individual security tools. In order to achieve its stated goals, security management must: (1) Manage the protection of data, applications, systems, and networks, both proactively and reactively; (2) respond to changes in business and organizational models as well as the changing threat environment; (3) integrate with IT infrastructure management and operations; and (4) all the while, maintain a level of security and operational risk that is pre-defined by that organization.

Identity Management

Identity management is the ability to identify every user, application, or device across organizations and provide flexible authentication, access control, and auditing while respecting privacy and regulatory controls. Delivered via a set of processes and tools for creating, maintaining, and terminating a digital identity, these tools allow administrators to manage large populations of users, applications, and systems quickly and easily. They allow selective assignment of roles and privileges, making it easier to comply with regulatory controls and contribute to privacy-sensitive access controls.

For HP, identity management is a pervasive set of technologies and solutions:

- Identity management is about the management of user, application, and device identities.
- Identity management is about the management of identities in different contexts: enterprises, SMBs, consumers, and the public sector.
- Identity management deals with the management of the entire lifecycle of identities and their attributes.

HP considers privacy management, identity services, business-driven identity management, identity-capable platforms, and device-based identity management as important emerging identity management fields and drives specific research in these areas from HP Labs.

As an example of an end-to-end identity management system, the HP National Identity Solution provides governments with a high-performance, extremely secure, and extremely reliable credentialing solution. Similarly, HP can provide fully integrated end-to-end identity management solutions to meet any enterprise or public sector need.

Trusted Infrastructure

As businesses and society increase their reliance on IT infrastructures, we face important challenges to stay ahead of security threats to infrastructure technologies. Fundamental IT building blocks must be innovated and redesigned to include security features. Across all technologies, from clients to servers, from networking to storage, and in printing systems, HP continually strives to improve infrastructure security mechanisms to support adaptive and flexible IT solutions.

HP is investing to ensure that we continue to deploy secure and reliable trusted infrastructures. HP is an industry leader, driving this agenda across platforms, operating systems, and infrastructure solutions. Importantly, HP's leadership in the Trusted Computing Group has brought the industry together, greatly increasing baseline security of infrastructure technologies to meet current and future customer needs.

Alongside other efforts, such as establishing secure development practices within HP and driving infrastructure technology standards, Trusted Computing provides the security building blocks that allow the IT industry to continue to innovate and deliver the power of IT across reliable trusted infrastructures.

HP Labs Invention

HP Labs security research contributes innovative technology breakthroughs across all aspects of the corporate security strategy. Focused research aligns directly with primary initiatives and drives business units to think differently about approaches to security challenges. From trust economics to trusted infrastructure and assurance automation, HP Labs is inventing new technologies for the full security lifecycle. In addition, HP Labs invests in longer-term research to sustain a competitive pipeline of invention and innovative security capabilities for a wide range of emerging technology and application domains.

Proactive Security for a Safer IT Environment

Today's enterprise environment is increasingly volatile due to changes driven by business opportunity and threats emerging from attacks that are ever more sophisticated. In addition, government regulation is increasing corporate accountability for proper business practices and for protecting individual privacy. These pressures mandate a change in tactics for IT security - a change to a new proactive approach rather than the conventional reactive approach.

To enable our customers to implement a proactive IT security environment, HP wants our products, services, and solutions to be secure throughout their lifecycle. By focusing on the key areas of proactive security management, identity management, and trusted infrastructure with keen attention to governance and compliance issues, we have developed a solid framework for proactive enterprise security. With this framework, we deliver a safer IT environment to our customers - one that responds to changes in threats and corporate business objectives while it maintains defined levels of security and risk.

Appendix A:
Principles of Design for Network Security

Appendix B:
Types of Firewalls and Open Systems Interconnection (OSI) Layers of Operation

Appendix C:
Authentication, Authorization and Auditing (AAA) Servers

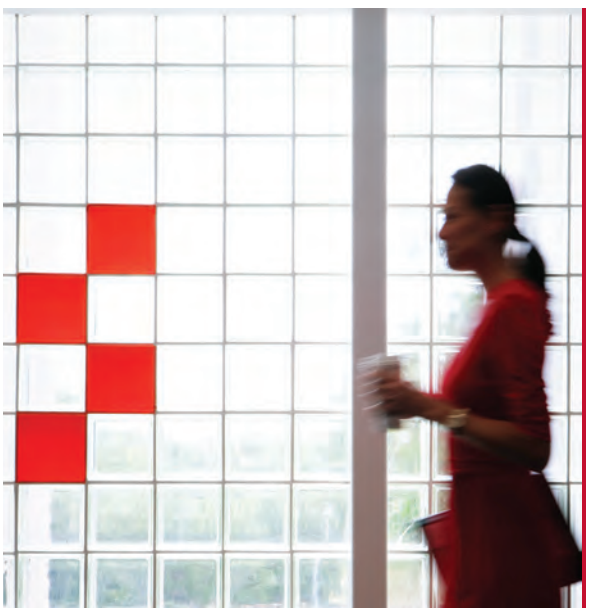

Appendix A

Principles of Design for Network Security

Standardization

Each type of network component, design, procedure, or baseline configuration has its own security implications. Each of these elements must be consistently managed and periodically reviewed as an organization evolves and its threat environment changes. Therefore, reducing the number of dissimilar elements in the network environment will, in general, reduce the complexity and cost of security.

For example, reducing the number of different operating system (OS) platforms reduces the number of job descriptions, operational procedures, administrative tools, and other supporting elements for which an organization must train users, identify threats, assess risks and vulnerabilities, and implement countermeasures. Furthermore, standardization of job descriptions, required training, and local team organization can significantly simplify security management.

Another advantage of standardization is the ability to deploy widely tested and trusted approaches throughout the enterprise. For example, standard protocols for secure communication such as Secure Sockets Layer (SSL), Secure Shell (SSH), and the IPsec (IP security) protocol family have been widely scrutinized and, over time, strengthened against a wide variety of potential attacks. By standardizing a limited number of well-accepted technical approaches and business best practices, organizations benefit from the experience and efforts of countless contributors over many years.

In some cases, however, implementation of diverse countermeasures (such as Linux-based bastion hosts to further secure a properly configured Microsoft Exchange e-mail infrastructure) can provide additional protection that outweighs the additional complexity. Therefore, the advantages of standardization should be balanced, in some instances, with the advantages of diverse countermeasures as part of a layered defense strategy.

Likewise, standardization brings with it an increasing requirement that whatever is standardized must be highly secure. A single vulnerability exploit can affect the entire network. Appropriate standardization can, however, conserve resources that can be applied to diverse countermeasures. For example, an organization could standardize on a regional e-mail infrastructure based on specific Microsoft Exchange and bastion host configurations.

Least Privilege Access

Individuals, systems, applications, and business processes should have access to the minimum amount of information necessary to conduct business. Least privilege access depends on the existence of a robust means of establishing and managing digital identities. (For more information about digital identities, see the Identity Management chapter.)

Least privilege access for networks has broad implications. It means that only public network resources should be available to individuals whose identity or organizational affiliation is unknown or unauthenticated. Individuals who must access the network, including individual network hosts, should only have access to the network-related information and equipment they need to do their jobs.

For example, operations personnel should not have access to application source code, and application developers should not have access to the resources needed for actively managing the production network. Physical access to data centers, telecom cabinets, and other network equipment should be restricted to authorized individuals.

Least privilege access involves more than people. An attacker could compromise any network resource. Resource privileges, and the privileges available through them, should be restricted to the minimum necessary to meet business requirements. For example, router access control lists (ACLs) and firewall configurations should be as restrictive as possible. Unnecessary services should be shut down on servers. Public resources such as web servers should be carefully secured to prevent unauthorized manipulation by external attackers.

Finally, least privilege access also involves the distribution of network traffic. IP networks are inherently insecure. Any wired or wireless network is a potential target for attackers seeking to observe or alter network traffic. On a LAN without switches, each datagram (message or message portion) reaches all host network interfaces, and it is up to the host to determine how the datagram is processed.

Encryption of sensitive information is an important countermeasure. All sensitive (non-public) information transmitted over wireless networks or the Internet should be encrypted. Encryption may also be applied to data stored on or transmitted over internal networks. However, encryption is not always feasible, and most encrypted network traffic remains susceptible to the analysis of communication patterns between network hosts, which is known as traffic analysis. In addition, if traffic from public network resources competes for bandwidth with internal network traffic, the organization is susceptible to a crippling Denial of Service (DoS) attack launched through the Internet.

In summary, traffic should not flow over a network segment unless there is a business need. Switches, which are commonly used to improve network performance, can also improve network security by channeling network traffic directly to its intended recipient or to a small subnet. Network designers can also use Virtual LANs (VLANs) - which define network membership on a logical basis rather than a physical basis - to resegment networks without rewiring them. Therefore, network designers can partition hosts and host traffic based on business and security requirements rather than physical location. Together, switches and VLANs can help prevent unnecessary distribution and exposure of network traffic.

Layered Defense

A layered defense is essential to an enterprise's network security strategy, approach, and implementations. Such an approach includes appropriate security policies, security awareness and training, security technology, best practices implementation and operation, and auditing.

Technology is critical to any layered defense strategy. An enterprise would not consciously connect internal networks to the Internet without a perimeter defense mechanism such as a firewall. However, even firewalls cannot be relied upon as the only way of protecting the network. Technology must be layered to provide maximum coverage and security for an enterprise's information assets.

The principal layers of security technology represent the perimeter, network, and hosts. Thus, in addition to a network-based Intrusion Detection System (IDS), a host-based IDS should be used to ensure host integrity. Encryption, anti-virus, system auditing and logging, backups, honeypots (hosts or other resources such as decoy user accounts used to lure and observe attackers), and other technologies support a layered defense strategy. Building a layered defense strategy requires breaking networks into divisions such as subnets and demilitarized zones (DMZs), with multiple layers of screening routers, firewalls, virtual private network (VPN) deployments, anti-virus solutions, intrusion prevention systems (IPSs), and IDSs to help identify malicious traffic not prevented by perimeter defenses.

Countermeasures must be combined to be effective. Any single countermeasure could fail or be susceptible to an attack, now or in the future. Similar to least privilege access, layered defense has broad applications to network security. For example, to log on to a system designed for internal use in a well-secured network, an attacker must penetrate multiple firewalls and routers secured with ACLs and somehow obtain a valid access credential. In addition, an IDS/IPS plays a role in mitigating the risk of unauthorized system access. If the attacker is a curious visitor who has obtained a valid user name and password by looking over an employee's shoulder, other controls must also be in place. These controls typically include site physical security, security policies, and awareness programs that shape employee behavior as well as logging and auditing of access to sensitive resources.

There are other important examples of the need for a layered defense. For instance, an employee downloading malicious software during an SSL browser session can circumvent many layers of classical network defenses such as firewalls and IDSs; consequently, additional defenses are necessary. These defenses typically include employee awareness efforts and policies against misuse of company resources as well as the enforced presence of personal firewalls and updated anti-virus software.

Network designers should think of a layered defense in three ways:
- Layers of different approaches that span physical, technical, and administrative controls
- Layers of physical and technical obstacles that a potential attacker must overcome
- Layers of countermeasures that prevent attacks, detect and report attacks, limit the damage that a single attack can carry out, and facilitate recovery from attacks

Redundancy

The network designer must consider the enterprise-wide impact of the failure or compromise of any network component. Redundant service providers, connections, entry points, and network services should all be considered. However, redundancy makes networks more complex and expensive. Therefore, it must be carefully justified. For example, redundancy may not be justified for a reliable switch that services a small workgroup, but Internet access for a major campus may well warrant redundant connections from different service providers.

Compartmentalization

Enterprise networks can be divided into compartments or subnets to control security and other operational risks, facilitate standardization, establish least privilege access, and achieve a degree of redundancy. Many organizations use the structure of their business operations as an initial guide to compartmentalization. For example, if an organization is partitioned into three major divisions and a corporate office, four business application compartments may be warranted.

One of the advantages of compartmentalization is that access policies can be determined centrally and implemented at compartment boundaries. In the example network of four business application compartments, hosts in each compartment may have limited access to hosts in other compartments because each division operates independently. However, hosts in the corporate network may have more extensive access to the divisional compartments to enable integration and oversight of divisional operations.

In addition to business application compartments, compartments can be created for other purposes. As shown in Figure A-1, compartments for cross-business services might be created by grouping e-mail, directories, and naming services; tools to monitor and manage the network; and all end-user desktops in separate compartments. Other compartments may contain applications that are accessible from the Internet (such as corporate and divisional websites) and hosts that are accessible to external organizations and individuals via secure remote access. Compartmental access policies generally allow most types of outbound traffic, but they limit inbound traffic based on business need or application type.

Figure A-1
Compartmentalized network, cross-business services

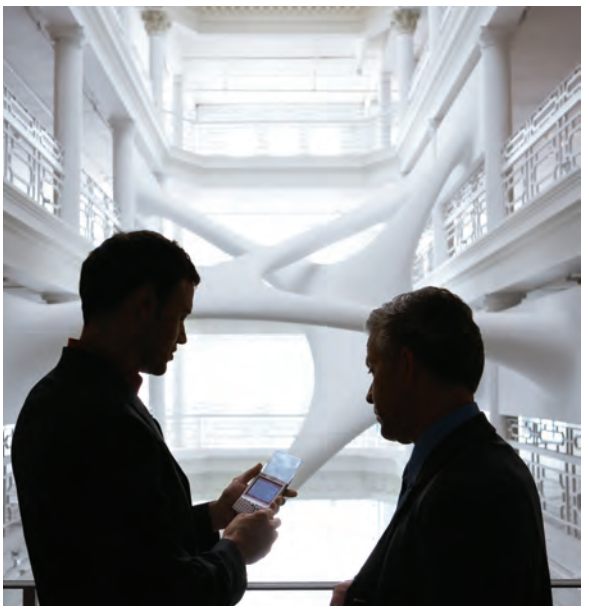

Compartments are not physical entities; they are accomplished by the logical network design. Hosts within a single compartment need not be in the same physical location. In fact, hosts within a single compartment can be located anywhere in the world, and diverse compartments may securely share a single site or computer room. A compartmentalized network can be engineered to adapt rapidly to changes in business structure and operations because compartments can be created, evolved, repurposed, or eliminated using VLANs with little or no change to the physical network. Compartmentalization facilitates adoption of other design principles. For example:

- Standardization is facilitated by centralizing policy management and providing standard network topologies for particular purposes.
- Least privilege access implementation is simplified by compartment access policies.
- Layered defense is facilitated by secured compartment boundaries that provide an additional layer of defense between the host and the network perimeter.
- Redundancy principles can be addressed by redundant resources implemented within the same compartment but at different physical locations. Multiple network routes may be established between compartment partitions located at different sites. Compartments themselves may be connected with a virtual backbone composed of redundant network routes.

Appendix B

Types of Firewalls and Open Systems Interconnection (OSI) Layers of Operation

Table B-1
Types of firewalls and OSI layers of operation

Type	*Layer of operation (OSI model)*
Packet filters	Network
Circuit-level gateways	Session
Stateful inspection firewalls	Network, transport, potentially others
Application proxy servers	Application

Packet Filters

Packet filtering is pervasive in today's network environment, and implementations exist in routers, switches, and OSs. Packet filters operate at the network layer and make decisions to allow or deny a particular network packet based on its content. Packet filters can be configured to allow or deny a packet based on the source or destination IP address; the User Datagram Protocol (UDP), Transmission Control Protocol (TCP), or Internet Control Message Protocol (ICMP) source or destination port; or the TCP acknowledgement bit. Packet filters are stateless, and they operate very efficiently due to the simplicity of their technique. However, they are vulnerable to spoofing (attacks based on falsified addresses and ports). In addition, they cannot defend against illogical packet sequences intended to disable or penetrate network hosts.

Circuit-level Gateways

Circuit-level gateways establish sessions between trusted hosts and clients. Like proxies, they enable clients and servers to communicate without a direct connection. Many circuit-level gateways are based on the SOCKS protocol, which enables clients that have been properly modified to use a SOCKS gateway to access TCP applications without revealing their IP addresses. SOCKS works with virtually any TCP application, including web browsers and FTP clients. SOCKS gateways can act as simple firewalls by examining incoming and outgoing packets and determining whether to allow them based on configured rules.

Stateful Inspection Firewalls

Stateful inspection firewalls allow or deny network traffic based not only on the contents of individual network packets but also on the state of existing conversations. This is crucial for preventing attacks that present unexpected or spoofed packet sequences to network hosts in hopes of penetrating them or denying service to others. In order to determine whether a particular packet is part of a legitimate interaction, these firewalls build and use state tables with information from all seven layers of the Open Source Interconnection (OSI) reference model. Stateful inspection firewalls are critical for protecting major networks.

Application Proxy Servers

Application proxy servers (proxies) are application-layer firewalls. Proxies provide security by breaking the direct connection between client and server, concealing network topology, and (in many cases) providing access control and communication security. Proxies add overhead by dividing each client-server connection into two connections, but they can also reduce network congestion by caching frequently used web pages. Proxy server software can run on dedicated or shared general-purpose systems, or it can be prepackaged as part of a proxy appliance. Some proxies mediate web access only; others handle a wide variety of protocols.

Forward proxies are placed in the client systems used to access the Internet an the Internet itself. Forward proxies can restrict Internet access, serving as one element of a layered defense against external attacks on internal systems. Forward proxies can be used to authenticate users and establish secure communication sessions with them. Reverse proxies are placed between Internet-facing applications and their users. Application users must communicate with the application through the proxy. Users may not be aware of this, since they use the application's domain name (for example, www.myapp.com) as they would with direct, non-proxied access. However, the Domain Name System (DNS) resolves the domain name to the IP address of the proxy rather than the application. Figure B-1 shows forward and reverse proxy configurations.

Figure B-1
Forward and reverse proxy configurations

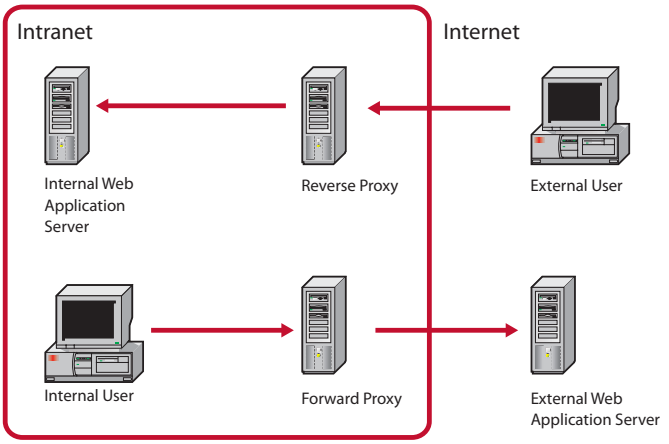

Firewall Network Architectures

There are three basic architectures of firewalls on networks: dual-homed host, screened host, and screened subnet. A dual-homed host (Figure B-2) has two NICs, each connected to a different network segment. The firewall controls traffic between the two networks. For example, in a very simple network, the firewall could allow selected outbound traffic from a subset of hosts and selected inbound traffic from the Internet to another group of hosts. Systems, like firewalls, that are properly secured against access from untrusted networks, such as the Internet, are called bastion hosts. Host security is discussed in the Trusted Infrastructure chapter (Chapter 4).

A screened host firewall (Figure B-3) is protected by a packet-filtering router that sits between the firewall and an untrusted network. The router's access control list (ACL) can be configured to allow traffic that meets specific source, destination, direction, port, and protocol criteria. Because the firewall receives only prescreened data, it can perform more detailed tests, such as stateful packet inspections, without adversely affecting network performance.

Although both dual-homed host and screened host firewalls provide basic security, most organizations require additional protection from attacks that originate on the Internet. The screened subnet (Figure B-4) network architecture includes a screened host firewall. It also segregates internal systems from systems attempting to access them from the Internet or another untrusted network. The Internet-accessible systems, configured as bastion hosts, are placed in a buffer zone or DMZ directly behind the screened host firewall. The internal systems are segregated from the Internet-accessible systems in their own subnet, which is separated from the DMZ by a router, a firewall, or both. In this way, a minimum of three devices must be compromised before an external attacker can reach an internal system.

Figure B-2
Dual-homed host

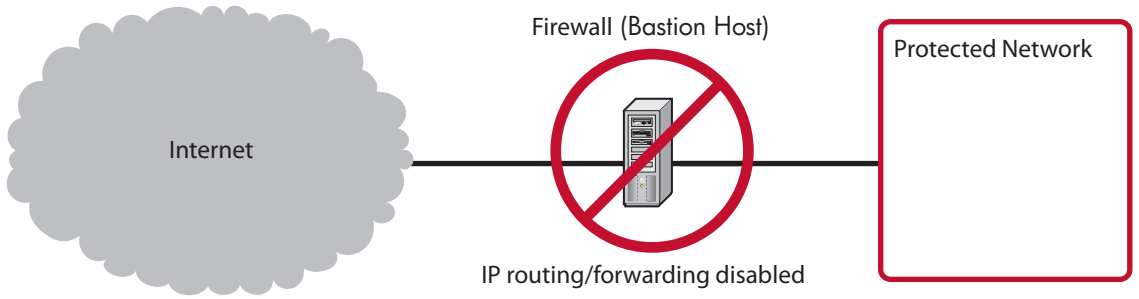

Figure B-3
Screened host firewall

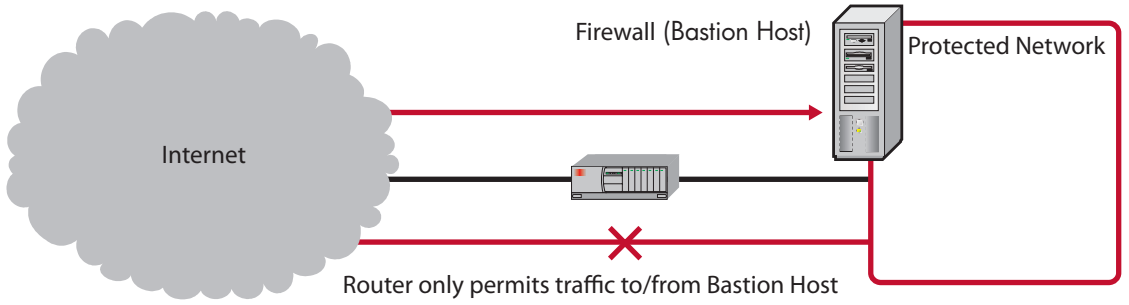

Figure B-4
Screened subnet

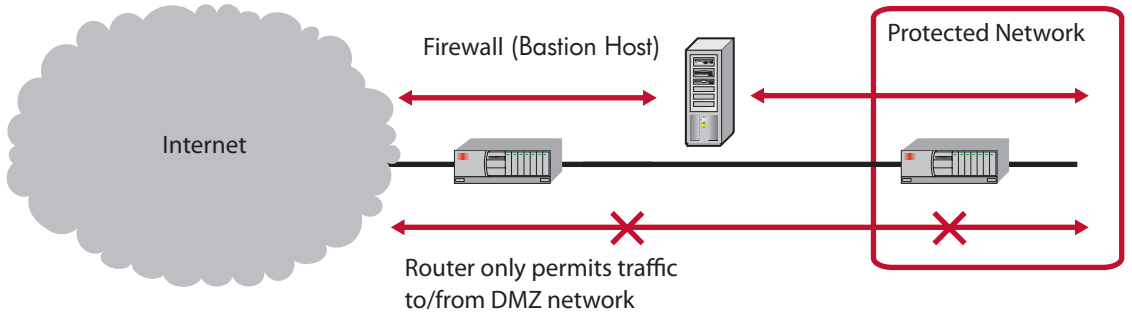

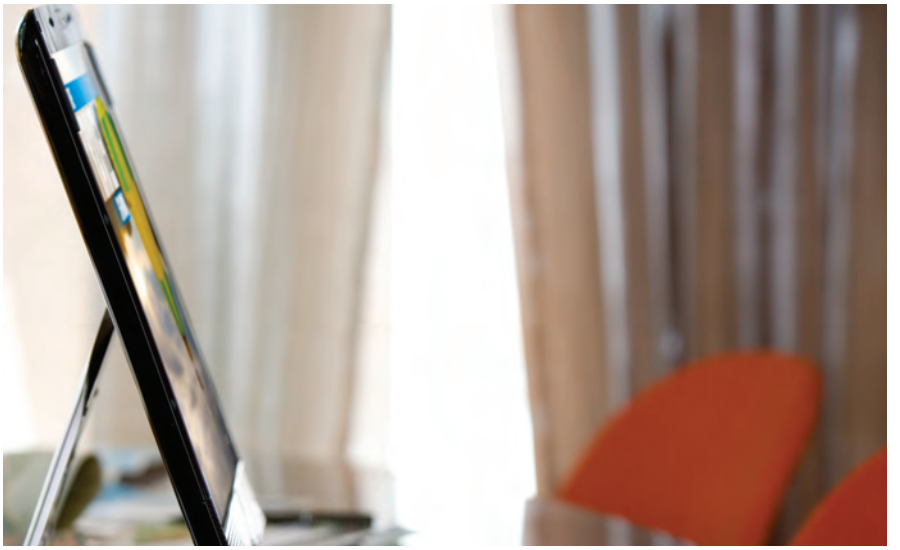

Appendix C

Authentication, Authorization and Auditing (AAA) Servers

AAA (Triple-A) servers authenticate network users, authorize them to use particular network resources, and audit their network usage. AAA servers provide a central control point for external network access, and they work with various types of network access servers that interact with users and collect their credentials. All AAA servers must support a client-server security model in which:

- The network access server collects users' credentials and requests authentication from the AAA server.
- The AAA server returns authorization information and other parameters.
- The network access server sets up a connection and writes an audit record.

AAA protocols must support multiple authentication methods, including user name and password, and multiple types of token authentication. They must also be extensible to accommodate future security requirements. There are three major AAA protocols today: RADIUS, Terminal Access Controller Access Control System+ (TACACS+), and DIAMETER (a play on the RADIUS acronym).

RADIUS

RADIUS, the most pervasive protocol, provides straightforward, efficient, and extensible services for authenticating individuals using a variety of credentials. It is available in a variety of implementations; however, it has no support for group membership, password management, account expiration, or event monitoring. Secondary authentication servers must be added to perform these functions in RADIUS environments. RADIUS uses User Datagram Protocol (UDP), which does not provide guaranteed delivery of messages between the network access server and the AAA server. It also does not, in its standard form, provide for confidentiality of client-server communication.

TACACS+

TACACS+ uses encrypted TCP packets for secure and reliable communication between clients and the AAA server, and it also logs system events such as access privilege changes. TACACS+ supports a wide range of security features compared to RADIUS, including group membership and privileges. TACACS+ can specify packet-filtering rules and access control lists (ACLs) for each session. However, TACACS+ is primarily a Cisco Systems protocol, and TACACS+ clients are principally Cisco appliances. In addition, its enhanced feature set and use of TCP give it greater network traffic overhead than RADIUS.

DIAMETER

DIAMETER is designed to overcome the limitations of RADIUS. It operates in a peer-to-peer mode; therefore, AAA servers can initiate requests themselves and handle transmission errors. It is also based on UDP, with enhancements for more reliable transport. Its other enhancements include support for roaming, cross-domain and brokered authentication, additional authenticable protocols, and enhanced security - including confidentiality and protection against replay attacks.

To learn more, visit **www.hp.com**

4AA1-7729EEW, March 2008

invent