SUN BLUEPRINTS
APPROVED

Developing a Security Policy

*By Joel Weise - SunPS℠ Global Security Practice and
Charles R. Martin - SunPS Java™ Centers*

Sun BluePrints™ OnLine - December 2001

http://www.sun.com/blueprints

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303 USA
650 960-1300 fax 650 969-9131

Part No.: 816-1953-10
Revision 1.0, 11/28/01
Edition: December 2001

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REPONDRE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.

Developing a Security Policy

A security policy is the essential basis on which an effective and comprehensive security program can be developed. This critical component of the overall security architecture, however, is often overlooked. A security policy is the primary way in which management's expectations for security are translated into specific, measurable, and testable goals and objectives. It is crucial to take a top down approach based on a well-stated policy in order to develop an effective security architecture. Conversely, if there isn't a security policy defining and communicating those decisions, then they will be made by the individuals building, installing, and maintaining computer systems; and this will result in a disparate and less than optimal security architecture being implemented.

This article discusses the importance of security policies for organizations that plan to use electronic commerce on the Internet; for government organizations that want to automate forms processing; and for any entity that may have external exposure of data processing environments. These organizations need some form of security architecture. This article also describes the basic steps through which security policies are developed and includes a set of recommended policy components.

In addition, this article is accompanied by a *Data Security Policy - Structure and Guidelines* template that was built on the recommendations made in this article. The template provides commentary; specific recommendations on all of the security topics chosen for the policy; and a detailed list of security policy principles. The template is available from:

`http://www.sun.com/blueprints/tools/samp_sec_pol.pdf`

The objectives of this article are to:

- Provide an overview of the necessity and criticality of security policies.
- Recommend a set of security policy principles that capture management's primary security objectives.
- Describe the basic characteristics of security policies.
- Describe a process for developing security policies.

Security Principles

The definition of security principles is an important first step in security policy development as they dictate the specific type and nature of security policies most applicable to one's environment. Security principles are used to define a foundation upon which security policies can be further defined. Organizations should evaluate and review these security principles before and after the development and elaboration of security policies. This will ensure that management's expectations for security and fundamental business requirements are satisfied during the development and management of the security policies.

The security policies developed must establish a consistent notion of what is and what is not permitted with respect to control of access to your data and processing resources. They must respond to the business, technical, legal, and regulatory environment in which your organization operates.

The principles here are based upon the following goals:

- Ensure the availability of data and processing resources.
- Provide assurance for the confidentiality and integrity of customer data and allow for the compartmentalization of risk for customers and your organization.
- Ensure the integrity of data processing operations and protect them from unauthorized use.
- Ensure the confidentiality of the customer's and your processed data, and prevent unauthorized disclosure or use.
- Ensure the integrity of the customer's and your processed data, and prevent the unauthorized and undetected modification, substitution, insertion, and deletion of that data.

Security Policy Fundamentals

This section provides basic information on the purpose, goal, definition, and implementation of a security policy. In addition, this section discusses the flexibility, communication, and management of an established security policy.

Purposes of a Security Policy

The primary purpose of a security policy is to inform users, staff, and managers of those essential requirements for protecting various assets including people, hardware, and software resources, and data assets. The policy should specify the mechanisms through which these requirements can be met.

Another purpose is to provide a baseline from which to acquire, configure, and audit computer systems and networks for compliance with the policy. This also allows for the subsequent development of operational procedures, the establishment of access control rules and various application, system, network, and physical controls and parameters.

Security Policy Goals

The goal of the security policy is to translate, clarify and communicate management's position on security as defined in high-level security principles. The security policies act as a bridge between these management objectives and specific security requirements.

Definition of a Security Policy

A security policy is a formal statement of the rules through which people are given access to an organization's technology, system and information assets. The security policy defines what business and security goals and objectives management desires, but not how these solutions are engineered and implemented.

A security policy should be economically feasible, understandable, realistic, consistent, procedurally tolerable, and also provide reasonable protection relative to the stated goals and objectives of management. Security policies define the overall security and risk control objectives that an organization endorses. The characteristics of good security policies are:

- They must be **implementable** through system administration procedures, publishing of acceptable use guidelines, or other appropriate methods.
- They must be **enforceable** with security tools, where appropriate, and with sanctions, where actual prevention is not technically feasible.
- They must clearly define the areas of **responsibility** for the users, administrators, and management.
- They must be **documented**, **distributed**, and **communicated**.

Policy Flexibility

A successful security policy must be flexible. In order for a security policy to be viable for the long term, a security policy should be independent of specific hardware and software decisions, as specific systems choices change rapidly.

In addition, the mechanisms for updating the policy should be clearly spelled out. This includes the process, the people involved, and the people who must sign-off on the changes.

Security Policy Communication

Once security policies have been established, they must be disseminated to all appropriate users, staff, management, vendors, third party processors, and support personnel. Given the nature of your enterprise, it may also be necessary to communicate some or all policies to customers as well. Establishing a record that those involved have read, understood, and agreed to abide by the policy is an essential part of this process.

Policy Management

To ensure that your policies do not become obsolete, you should implement a regular review process of them. That process should include some form of update mechanism so that changes in your organization's operating environment can be quickly translated into your security policy.

We recommend that a specific organization be identified, such as the data security department, and be chartered with the custodianship of your security policy. That organization would then be responsible for conducting a regular review and as applicable, updating your security policy.

Relationship to Standards and Procedures

Security policies embody management's overall security expectations, goals and objectives. To be practical and implementable, policies must be further defined by standards, guidelines, and procedures. These must ensure that all operations are consistent with the intent of the security policies.

Standards, guidelines, and procedures provide specific interpretation of policies and instruct users, customers, technicians, management, and others on how to implement the policies. Your organization should undertake the definition of standards, guidelines, and procedures only after the development and acceptance of security policies, and after specific security mechanisms supporting these policies are determined or implemented.

Implementation in IT Systems

Once policies, guidelines, and procedures are established, appropriate parts of these procedures can be implemented in information-processing systems. In Java™ technology-based systems, for example, some of the policies might become requirements for third-party software such as cryptographic providers, while others would become definitions of roles and access rights to be implemented using the Java Security APIs.

It is important to emphasize, however, that implementation of some policies in information processing systems is *never* sufficient to ensure that the customer’s trust is maintained: It is *at least* as important that all policies and procedures external to the information-processing system are also correctly enforced and performed. Without correct operation, security can never be guaranteed.

Security Stance

Security policies should be designed to provide effective and economically efficient directives for risk mitigation. Policies should be based upon a formal security stance that is determined by management. In general, a security stance of 'least privilege' can be used in most commercial, financial, Internet, and governmental environments. A typical security stance is:

> “All data defined as confidential must be protected on a need to know basis only to properly identified and authenticated entities, in all of its forms and on all media, during all phases of its life, from generation to destruction, such that it cannot be compromised, released to any unauthorized entity, or otherwise have its confidentiality or integrity placed at risk.
>
> All processing resources, including all applications, systems, network, hardware and software, are only accessible on a need to know basis, only to properly identified and authenticated entities.”

Security Policy Structure

The basic structure of a security policy should contain the following components:

- A statement of the issue that policy addresses.

- A statement about your position on the policy.
- How the policy applies in the environment.
- The roles and responsibilities of those affected by the policy.
- What level of compliance to the policy is necessary.
- What actions, activities and processes are allowed and which are not.
- What are the consequences of non-compliance.

A list of specific topics is also noted in the *Sample Data Security Policy and Guidelines* template listed in the *References* section.

Roles and Responsibilities

The development of security policies is predicated upon the participation of various organizations. In general, it is recommended that the following areas participate in this development effort:

- Business management
- Technical management
- Data security
- Risk management
- Systems operations
- Application development
- Network engineering
- Systems administration
- Internal audit
- Legal
- Human resources

Audience

In addition to determining the roles and responsibilities of those involved in the development of security policies, the intended audience for your security policies must also be considered. In general, you must first determine if policies are internal or external or both; and then determine the orientation of the audiences by topical area such as:

- Customers or clients
- Executive management
- Business management

- Technical management
- Employees, temps, contractors

Recommended Development Method

We recommend a policy development approach that is consistent with industry best practice as outlined in current standards such as ISO 17799. In general, this approach uses a risk assessment of one's business and technical environment to establish the tolerance for risk and understand minimum or baseline security requirements. This risk analysis enables the identification of threats and countermeasures, and the subsequent development of specifically tailored policy statements.

The following provides an outline of the tasks used to develop security policies.

1. All responsible organizations and stakeholders are identified and their roles, obligations and tasks detailed.

   It is important to understand how your organization is structured, who will be the responsible owner of the security policy and also who will function as its custodian. Most importantly, it is critical to obtain the appropriate level of consensus to ensure that the security policy properly reflects the issues, concerns, requirements, goals, and objectives for your organization. Representation should be as broad as practical but at a minimum include: data security, legal, human resources, internal audit, operations, and development organizations.

2. The primary business objectives are outlined.

   Knowing the primary objectives of your business is important to scoping the security policy effort. For example, one organization may require extensive audit, monitoring, and backup and recovery processes because of regulatory mandates while this may not be applicable to another. The intent here is to make security policy cost effective. That is, do what is appropriate for your organization, not the security consultant selling you the security policies.

3. A list of security principles representing management's security goals is outlined.

   Accompanying this article is a list of security principles. These should be reviewed and incorporated into your security policy development effort as necessary. The purpose of the security principles is to allow your organization to state in a plain and simple fashion, without technical details or jargon, what core values are most important to your organization.

4. All applicable data and processing resources are identified and classified.

The method that we recommend for security policy development uses a data-centric model. In today's IT environments, data is often one of the most important assets and should be treated accordingly. For that reason, cataloging your data and processing resources enables you to more easily make qualified and informed decisions about their use and value. This then enables you to later apply the most cost effective controls on those assets.

5. A data flow analysis is performed for the primary data classifications, from generation through deletion.

   As noted above, we recommend a data-centric model for policy development. The purpose of a data flow analysis is to allow you to identify all of the trust points that touch your data. For instance, in a transaction processing system, data may flow through browsers, web, data, and other servers or firewalls and be stored in databases, on magnetic tape or paper. By tracing the flow of your data assets through your processing assets, you can later determine the type and placement of logical and physical controls to protect those assets.

6. The primary threats that can reasonably be expected in one’s environment are outlined.

   The development of a threat profile enables you to decide what type of threats exist in your particular environment, what the probability is of a threat manifesting itself into an actual problem, and what the ramifications, costs and consequences are of those threats being realized. Remember, threats vary widely between different environments. The threats and consequences of attacks to a financial network processing monetary instruments will be different than the threats and consequences of attacks to an online photo gallery that only displays art.

7. The primary security services necessary in the environment are identified.

   After your data and processing assets are identified and a threat profile created, the next step is to determine what general security services would be appropriate in your environment. These security services are high-level and can include for example: accountability, authorization, availability, identification, authentication, confidentiality, integrity, and non-repudiation. Knowing what security services your environment requires will drive the selection of the types of security policies you will need as well as the specific content or components of those policies.

8. A generic policy template is constructed.

   The structure of a security policy can take many forms. This article offers recommendations for both the components and characteristics of security policies. This step is used to articulate the specific topics that you consider necessary for each security policy. Refer to the *Sample Data Security Policy and Guidelines* template listed in the *References* section for some examples of recommended topics.

9. A list of security policies is defined.

   The last step before actually drafting the security policies themselves is to identify all of the security policy focus areas that must be addressed. The creating of this list is based upon the results of the above steps. A sample list of recommended focus areas is included in the *Sample Data Security Policy and Guidelines* template listed in the *References* section.

Assumptions

A discussion on the development of security policies requires making certain assumptions. This article presents many data security topics that are presented as examples and should be treated accordingly. Not all subjects are applicable in all circumstances. They are provided to ensure you receive a sufficient body of best practices to enable the development of a comprehensive set of security policies.

Security Concepts

Some fundamental security concepts that should be considered when undertaking policy development are described in this section. These concepts are provided as background material to enable you to properly scope your policy development effort.

Trust

Underlying all security policy, procedures, and architecture is the expectation that the policy, procedures, and architecture will preserve as confidential that which should remain confidential. In other words, you want to *trust* the system.

This notion of trust is the ground on which all security stands. In order to create this trust, you must understand and agree to a security policy, and must have confidence that this policy will fulfill their expectations.

Risk

In the past, the most common security policy has been *no* security policy: security was left up to individual implementors, and often overlooked. Following the first damaging security breach, the most common second policy is one so restrictive that it is ignored.

In order to construct a security policy that will neither be overlooked, nor ignored, it is necessary to make certain the security policy reflects realistic business goals and business values. This is done most easily by taking a *risk-driven* approach.

In finance, risk (R) is defined as:

$$R = H \times P_e$$

In this equation, H is the *hazard*, or the cost of the undesired event, and P_e is the *probability* of the undesired event. Thus if the cost of an undesired event is large, but the probability of that event is very small, the actual risk is still small.

Consider, for example, the hardened auxiliary power supply for a banking data center. The power supply itself costs $10 million; thus the *hazard* involved in the loss of the power supply is $10 million (plus whatever opportunity costs might be involved, which for simplicity we ignore here). We estimate the probability of a complete and catastrophic failure of the power supply—say by destruction of the power supply by terrorists—as one in a billion, or 1×10^{-9}. Thus the risk is $10 million × ($1 \times 10^{-9}$), or $0.01—one cent.

On the other hand, consider the personal checking accounts at the same bank: if the bank uses a 4 digit PIN number as the sole means of authentication, then the probability of guessing the PIN correctly in one try (and thereby compromising the account) is 1×10^{-4}. If the average checking account balance is $3000, then the risk of the checking account being compromised is $3000 × ($1 \times 10^{-4}$), or $0.30. Surprisingly, the risk to the bank of one checking account being compromised is *thirty times* the risk of the $10 million dollar power supply being lost.

An easy way to understand this is by considering the cost to the bank of insurance to cover each hazard. An insurance company, in essence, bets its customers that the undesired event will happen rarely enough to allow the insurance company to make a profit on its premium. Since the risk of a checking account being compromised is about 30 cents, this premium has to be slightly more than 30 cents per account.

In general, of course, we don't quantitatively analyze the risk being managed by every specific element of the security policy. An understanding of the nature of risk and how it is quantified, however, can help determine whether any specific element of a security policy is appropriate and suited to the business goals.

Holistic Security

It is essential to realize that *no one component of a system provides security.* Security policies are only effective in the context of an integrated and comprehensive data security architecture; that is, all access control, firewall, cryptographic, key management, and similar mechanisms must work in a coordinated, cohesive and coherent fashion. Thus, in order to provide security, it is not sufficient to simply implement some security features in an IT system. Real security requires us to holistically integrate risk management controls including people, processes, and technology together; security policies must reflect this holistic approach.

Data Centric Approach

The development of security policies is based upon a data centric model. We consider the data of an IT system to be the primary asset at risk. This means that the type of data, its ownership and classification are considered, as well as the flow of data through your organization, systems and networks. Following the data through the system is an effective way to develop efficient and meaningful security policies.

Conclusion

This article has discussed the most important, and often least understood, aspect of security: the security *policy*. A security policy establishes the expectations of the customer or user, including what their requirements are for confidentiality, integrity, and appropriate management of their data, and the conditions under which they can trust that their expectations are met.

A security policy does not, in itself, establish the requirements of a customer on specific information systems. It is instead the bridge between the customer's expectations, and stated requirements that can be applied to develop an information system.

A security policy should clearly state the customer’s expectations, and should be based on an evaluation of the risk to a customer should the customer’s expectations not be met. This risk-based evaluation helps avoid an infeasible, intractable, or excessively restrictive security policy.

In order to ensure the policy correctly describes the expectations of all stakeholders, this article is accompanied by a template available from the Sun BluePrints™ Web site (`http://www.sun.com/blueprints/tools/samp_sec_pol.pdf`) which describes an outline business process for development of a security policy.

Additionally, to simplify the statement of a complete and effective security policy, the template accompanying this article also includes an outline of the necessary components of a security policy, and discusses the appropriate contents for each component. If applied with care and thought, this template should allow a well-documented security policy to be developed.

The security policy is the foundation on which effective security is built. As with any foundation, it must be well designed, and well constructed; it can then be trusted to support the customer's needs effectively, and enduringly.

References

- Weise, Joel and Martin, Charles R., *Sample Data Security Policy and Guidelines Template*, Sun BluePrints OnLine, December 2001, `http://www.sun.com/blueprints/tools/samp_sec_pol.pdf`
- `http://web.mit.edu/security/www/GASSP/GASSP.DOC`
- `http://www.faqs.org/rfcs/rfc2196.html`
- `http://www.faqs.org/rfcs/rc2504.html`
- `http://csrc.nist.gov/isptg/html`

Author's Bio: Joel Weise

Joel Weise has worked in the field of data security for over 20 years. As a Senior Security Architect for Sun Professional Services, he designs system and application security solutions for a range of different enterprises from financial institutions to government agencies. He specializes in cryptography and public key infrastructures. Prior to joining Sun, Joel was a Senior Project Manager for Visa International. There he was responsible for developing cryptographic standards, designing key management and cryptographic systems and architecting security solutions for chipcard, Internet, and other new products.

Author's Bio: Charles R. Martin

Charles R. Martin has been in the computer business for more than 30 years, and involved with computer security since 1983. He was the original architect for a DARPA B3/A1 X server, developed methods for covert channel analysis in full-scale UNIX® kernels, and has co-authored several chapters in the NRL handbook on trusted system evaluation. He is currently a Senior Java Architect with Sun Microsystems in the Sun Java Center developing methods for quantitative analysis of distributed system architectures, and a member of the graduate faculty at the University of Colorado in Boulder.