

Essay 7

Information Security Policy

Ingrid M. Olson and Marshall D. Abrams

This essay discusses information *security* policy, focusing on information control and dissemination, for automated information systems (AISs). Most organizations have some sort of high-level information policy that addresses how and what information is to be handled by the organization. AISs have changed how information can be used. A further refinement of the high-level information policy is necessary to deal with this automation and establish what is considered acceptable behavior with respect to the information. This refinement process involves determining the appropriate set of policy-oriented limitations. It can take place at many levels, from a top-level corporate decision to a hardware implementation choice.

An information security policy addresses many issues such as the following: disclosure, integrity, and availability concerns; who may access what information in what manner; basis on which the access decision is made (for example, user characteristic such as nationality or group affinity, or some external condition such as time or status); maximized sharing versus least privilege; separation of duties; who controls and who owns the information; and authority issues. In the past, R&D has focused primarily on DoD policies based on user clearances and data classification, but many other access control policies are in use in the manual world, in other government agencies, and in the private sector. Policies such as a press release policy (sensitive until released at <time, date>), access based on roles (only vice presidents and above have access), and many others are real and useful policies with special characteristics not easily handled by most current systems. This essay discusses some of the aspects that must be considered when developing an information security policy for a given organization.

A policy is a plan or course of action, designed to influence and determine decisions, actions, and other matters [AMER82]. Organizations typically have many policies governing all aspects of their operations, security being one major consideration. The Information Technology Security Evaluation Criteria (ITSEC) define a corporate security policy as follows [COMM91]:

The set of laws, rules, and practices that regulate how assets including sensitive information are managed, protected, and distributed within a user organization.

An organizational security policy has been defined as follows [STER91]:

The set of laws, rules, and practices that regulate how an organization manages, protects, and distributes resources to achieve specified security policy objectives. These laws, rules, and practices must identify criteria for according individuals authority, and may specify conditions under which individuals are permitted to exercise their authority. To be meaningful, these laws, rules, and practices must provide individuals reasonable ability to determine whether their actions violate or comply with the policy.

This essay will focus on a subset of the organizational security policy: - the information security policy governing the protection of *information*. Information is one of many resources an organization must protect. While the protection of information has always been a major concern for organizations, the computerized automation of information has vastly changed the threats to and vulnerabilities of the information, resulting in a need to further interpret and refine the information security policy for the automated information system (AIS) environment.

The ITSEC has defined a technical security policy as follows [COMM91]:

The set of laws, rules, and practices regulating the processing of sensitive information and the use of resources by the hardware and software of an IT [information technology] system or product.

There are many aspects to a technical security policy or an information security policy. Areas such as the labeling of the information, accountability, information ownership, modification of the information, audit, and dissemination controls must all be addressed for the policy to be complete. This essay focuses primarily on the information dissemination/control aspects of the policy — that is, who can access information within the system and what they can do as a consequence of such ac-

cess. In addition, the discussion focuses on the AIS-related aspects of the policy. Other implementations of the policy may be possible through personnel, physical security, or procedural controls.

Policy refinement

At a high level, the information protection objectives of many organizations look very similar. For example, the policy may state “sensitive information processed by the organization’s resources shall be properly safeguarded against accidental or malicious disclosure, alteration, destruction, or delay” [OLSO90]. For different organizations, though, this statement can have very different implications. Within the government, this sensitive information may be designated sensitive unclassified and be subject to the Privacy Act and/or the Computer Security Act. National-security-related sensitive information is classified according to complex rules and is subject to Executive Order 12356 and numerous Department of Defense directives. Private organizations may identify sensitive information as proprietary, personnel confidential, source selection sensitive, and so on, each with its own implications of who may access the information and under what conditions.

The process of refining an organization’s high-level policy into an implementable AIS security policy involves many choices and decisions, from top-level decisions concerning organization objectives to hardware implementation choices. Through this refinement process, there will be many representations of the policy. At the higher levels, the policy is likely to be written in a natural language, which is easy to understand in a general way but is also subject to ambiguity. At some point, a precise formal language restatement in a formal security policy model may be appropriate. The essay “Formal Methods and Models” (Essay 8), by Williams and Abrams, discusses these different representations in more detail.

National policy, other organization policy, international standards, particular project requirements, political issues, implementation platform limitations, or other factors may influence or determine some of the policy refinement decisions. Some of the choices may have legal, ethical, or privacy issues driving the decision. For example, should the stored record of a legal transaction be in its original form or is another format acceptable? Will the organization monitor the activities of its employees? Does the environment foster open exchange of information and allow access to everyone unless specifically denied, or is information closely held and access denied to anyone unless specifically authorized access?

At the implementation end, limitations of existing systems may not allow for precise implementations of the defined organization policy. In the past, research and development have focused primarily on Depart-

ment of Defense (DoD) policies based on user clearances and data classification. Some security mechanisms developed to support these DoD policies may not be useful for supporting the many other information security policies in use in other government agencies and the private sector. Therefore, other nontechnical (for example, administrative) controls may need to be implemented around the AIS to support the policy objectives, or the organization may accept certain security compromises in exchange for the benefits of automation. So, while a high-level policy decision requires mechanisms available to support the decision, hardware limitations may also restrict some of the possible choices. The overall architecture of the system may drive some of the policy choices as well. Decisions made for a stand-alone system may be quite different from those made for a large central computer room facility supporting hundreds of users over a local area network, or a distributed system spread out over a campus of buildings, or a cross-country communication network.

The rest of this essay discusses some of these refinement decisions that must be made in developing an implementable information security policy and some of the implications of the decisions.

Information security objectives. Information security or information technology (IT) security has long been considered to consist of three main objectives: the preservation of the information's confidentiality, integrity, and availability [COMM91]. Donn Parker, a founding father in the field of information security, has also suggested adding to the list of objectives authenticity ("the valid representation of that which it is intended to represent") and utility ("the state of being useful or fit for some purpose and designed for use or performing a service") [PARK91].

One of the first decisions in the policy refinement process is the prioritization of these security objectives. In many cases, the written policy will emphasize only one objective; however, for most environments all objectives are of some concern.

Confidentiality, or the prevention of unauthorized disclosure of the information, has been the primary security objective of many of the DoD security efforts. The Trusted Computer Security Evaluation Criteria (TCSEC) have confidentiality as their primary concern. One normally thinks of the actual information as being the concern for protection. In a communication system, though, in some cases, knowing who is communicating can be as sensitive as what is being communicated. The frequency and volume of communication could also be very sensitive. For example, increased message traffic to a military base could indicate an upcoming activity, or increased traffic between two financial entities could imply a financial strategy [SIMP90]. Methods of disclosure may also

vary, and further refinements of the policy should address all applicable areas. Unauthorized disclosure may be accomplished through

- wiretapping or eavesdropping,
- unauthorized access to the information in the AIS either by unauthorized users (for example, hackers) or by authorized users accessing information to which they are not authorized,
- printouts of sensitive information sent to unattended printers in public areas, or
- large amounts of information leaving the organization on a floppy disk.

Having a detailed policy about required technical controls to forestall hackers will not adequately protect the organization if other aspects of the policy do not prohibit authorized users from wrongdoing.

Integrity, another information security objective, is a current research topic. There is no accepted single definition of what integrity encompasses. A multipart definition of integrity has been formulated as follows [NCSC91a]:

1. A subgoal of computer security which pertains to ensuring that data continues to be a proper representation of information, and that information processes continue to perform correct processing operations.
2. A subgoal of computer security which pertains to ensuring that information retains its original level of accuracy.
3. Sound, unimpaired, or perfect condition [NCSC88].

Two examples of the areas to consider in the policy refinement process are (1) in a communication system, changing the order in which messages are received or when they are received, and (2) repudiation: the denial of either the receipt or the origin of a communication.

The availability objective is generally seen as ensuring that the system is available to authorized users when needed. In life-support systems, for example, the availability objective is paramount. The recent widely published incidents of the Internet worm and the large-scale telephone outages in major metropolitan areas demonstrate the implications of denial of service. To date, little research has been done on availability concerns. The Canadian government has recently issued some guidance in this area.

Granularity of controls. A policy may affect few or many users. Some policies are intended to apply to all users — for example, on a DoD system, the policy enforcing the concepts of users' clearances and data classifications. Other policies may apply only to a specific application or

type of information. The coarsest granularity is concerned with access to the entire computer system. If there is access to the entire system, when users are authorized access to the system, they have access to everything on the system. For single-user systems or systems dedicated to a specific task, this granularity of control may be sufficient. However, for multiuser systems running several applications, a finer granularity of control is probably required. Performance and overhead considerations are key in determining the level of granularity. The flexibility of the policy and the ability of the supporting mechanism to be tailored as needed are also considerations.

Another aspect of granularity is the issue of centralized versus decentralized control. With centralized control, a single authority controls all security aspects of the system, reducing the complexity of the security controls and the administrative demands on other users, while creating a potential single point of failure and bottleneck. Decentralized control, on the other hand, while technically and administratively more complex, puts the responsibility for many security functions in the hands of individuals, users who are probably most familiar with the particular requirements of the information.

Authority. A vital part of defining the information security policy is defining the authority for the policies and providing for the delegation of authority. The strength, scope, and span of the AIS controls depend, at least in part, on the authority of the person or organization that makes the rules and maintains the information and rules used by the system.

At the highest level, for example, the country's president and national policy or the chairman of the board and corporate policy are the top levels of authority for the policy. The level of authority is delegated as the policy is refined, and at some point, we reach the boundary between the administrative controls among people and the technical controls within the computer. Defining this hierarchy of controls is part of determining authority, and even within the AIS there will be levels of authority. For example, a user may have authority over his own working files, the project manager has authority over the project group's files, and the system administrator has authority over system files. In addition, there may be a security officer who has authority over all files.

Another aspect of defining levels of authority in a system involves clearly defining the different types of users and the responsibilities and authority of each. Many systems support the concept of groups and roles. The policy should also address related issues such as who defines group membership, when the use of groups is appropriate, whether a user can belong to more than one group, what the individual accountability requirements are within the groups, and how to resolve conflicts between individual user and group privileges. As an example, a system may define four types of users:

1. *User*. One who has authorized access to information on a computer. Authorizations may include the ability to read, write, delete, append, execute, and grant/rescind permissions to some objects. For example, the user may be granted all permissions for his own working files, read-only access to project files, and no access to system files. For some information, the user may be the owner or custodian.
2. *Owner*. That individual manager or representative of management who has the responsibility for making and communicating judgments and decisions on behalf of the organization with regard to the use, identification, classification, and protection of a specific information asset. For example, the owner of the information may be the only one authorized to grant/rescind user privileges to access the information.
3. *Custodian*. One having authorized possession of the information and entrusted by the owner to provide proper protection in an ongoing operational environment.
4. *Security administrator*. The person responsible for the security of a system. Functions that the security administrator is expected to perform include auditing, initializing, and maintaining the security parameters of the system.

Policy decision attributes. Another aspect of refining the information dissemination/control policy is determining what information needs to be maintained about the users and the information being accessed to make the access control decision. The inputs to access control decisions are attributes of the user (for example, identity or clearance level), attributes of the information being accessed (classification level, document number, source of information), or some attribute of the environment or context of the system (time of day, status). Each system must select the relevant information to make an access control decision. In general, a policy specifies a comparison of subject and object attributes and/or context, but some policies may involve only one of these sets. Some attributes may support many policies (for example, user identity), while others have a one-to-one relationship with the policy (for example, source of the document may be used to support an acquisition support system). The following paragraphs briefly describe some of the attributes more commonly used. They are divided into five main categories: user characteristics, object characteristics, external condition, data content versus context attributes, and others.

User characteristics. User characteristics are commonly used in an access control decision process. Conceivably any attribute of a user could be used (age, sex, residence, place of birth, and so on), as identified by

the appropriate policy. The following are some attributes used in current policies:

- *User clearance level.* This attribute is based on national policy that requires the protection of sensitive national-security-related information based on a user clearance level and the information classification. Traditional mandatory access controls have been developed to handle access control based on this attribute. However, there is little consensus on the applicability of these types of controls outside of systems handling the hierarchical user clearance/information classification scheme.
- *Need-to-know.* An attribute may be associated with a user that indicates that the user has the “need-to-know” for a certain type of information. This type of attribute is commonly used in the intelligence community. It also may apply in other sensitive applications, such as payroll, where only the payroll clerk and an employee’s supervisor have the “need-to-know” for an employee’s salary information.
- *Role.* Associated with a user may be the role in which the user is currently acting. For example, there may be the role of “ordinary user,” which does not allow for any special privileges. Other role examples include position title, place in the organization, or function, such as security officer, data entry clerk, or department manager.
- *Group affinity.* This attribute might include nationality, employer, or user organization. Within the Department of Defense, the designation NOFORN (not releasable to foreign nationals) uses the nationality attribute, as does the marking NATO (which implies individuals from certain countries may have access if they have the appropriate need-to-know). Exceptions to either of these policies may also be specified in the form Releasable to X. NOCONTRACT (no contractor) is a policy used in the federal government to indicate only government employees may have access. For this policy, the employer attribute could be used. Company Confidential is a similar policy used in the private sector to limit access to individuals within the organization. Today, most of these policies are handled administratively.

Object characteristics. Other common attributes used in the access control decision process are attributes associated with the information being accessed. Several examples of these attributes follow:

- *Sensitivity labels.* Within the Department of Defense, there is a well-defined structure for labeling information with a hierarchical

classification level (Unclassified, Confidential, Secret, or Top Secret). In addition, within the DoD/intelligence community, there are numerous compartments, categories, handling restrictions, and other markings used to further restrict access to information, and well-defined policies stating the access rules — for example, LIMDIS (limited distribution), ORCON (originator controlled), and PROPIN (proprietary information).

- *Information identifiers.* Other identifying attributes of the information could also be used — for example, the source or originator of the information, the owner of the information, and the document number.
- *Access control list.* Associated with the information may be a list of who is authorized to access the information. The owner of the information might be allowed to specify the list according to whatever criteria he chooses — for example, “who do I like” or need-to-know.

External condition. Some policies might be based on attributes associated with some external condition (context), such as time, location, or status. In addition, most of the other attributes previously discussed were assumed to be relatively static. Some policy decisions may be based on data that is expected to change:

- *Location.* Access may be based on location (for example, only a user from the main office is allowed access to some information).
- *Time.* Access to the information may vary with time (a press release policy where the information is highly sensitive until 9:00 a.m. on Tuesday morning, when it is made public).
- *Status.* A status variable may be maintained that reflects some condition in the real world that affects the policy decision made (crisis or exercise status).

Data content versus context. Some access control policies may depend on the value of the data. For example, user *X* may not be allowed to see the personnel file of anyone earning more than \$20,000. More complex policies may depend on context — that is, the identity of other data fields. The use of static labels is a carryover from manual methods of determining the sensitivity of information. The dynamic determination of access based on content and context has been recognized as a potential replacement for static labels and is currently a topic of research, particularly in the database area [SMIT88]. The rules for such dynamic real-time evaluation of content and context are very complex and probably not well understood in the manual world.

Others. Most automated computer security and access control today focus first on secrecy policies. As research continues and technology develops to support the integrity and availability **Error! Bookmark not defined.** concerns, other types of attributes will surely be needed to support those areas as well.

Summary

This essay has reviewed some of the policy decisions necessary to refine a high-level information dissemination/control policy into an implementable access control policy. Information security policy issues that were discussed with regard to policy refinement include information security objectives, granularity of controls, authority, and policy decision attributes.

Conclusions

Many types of information security policies are in use today in the private sector and the federal government, but technology still has a long way to go to adequately and efficiently support most of these policies. However, the process of refining an organization policy is critical to understanding the security requirements of the information, not only to adequately protect the resources of the organization, but also to guide the direction of technology to provide the automated support for these policies.

