Chapter 1
Introduction

The Information Assurance Technical Framework (IATF) exists to address questions such as:

- How do I go about defining information protection needs and solutions?
- What technology exists to give the protection I need?
- What organizational resources are available to help locate the protection I need?
- What kind of markets exist for Information Assurance (IA) products and services?
- Where should research in IA approaches and technology be focused?
- What are the principles of IA?

This evolving document is published to provide recommendations and information on current information assurance concerns and practices to System Security Engineers and others who address IA in their work. Over time it will reflect changes in policy, technology, environments, and the uses made of systems that depend upon information.

1.1 Objectives

The Framework has several objectives:

- Raise the awareness among users of information-dependent systems of information assurance technologies.
- Identify technical solutions to IA needs in accordance with national policies.
- Employ the technology focus areas of a *Defense-in-Depth* strategy to define approaches to IA.
- Define the security functions and protection levels needed for different situations or mission scenarios (referred to as “cases”).
- Present the IA needs of users of information-based systems.
- Highlight the need to engage a team of IA or information systems security experts to resolve pressing security needs.
- Aid the development of IA solutions that satisfy IA needs by highlighting gaps in the currently available commercial and government protection technologies.
- Provide guidance for solving IA issues by offering tutorials on available technologies, trade-offs among available solutions (at a technology versus product level), and descriptions of desirable solutions’ characteristics.
- Assist purchasers of IA products by identifying important security-related features that should be sought.

1.2 Intended Audiences

The Framework addresses the needs of several groups of people. The following describes each group and indicates how the document can be used.

- **System security engineers.** To assist in developing IA solutions tailored to a particular customer's needs. The customer's needs can be compared with the various Framework technology areas, cases, and recommended solutions. From these, a tailored solution can be created for this particular customer.
- **Customers.** To provide answers to the myriad issues and technical challenges involved in selecting adequate IA features and assurance levels for their system and networks. Customers can include system users, managers, and security officers or administrators. With this knowledge, customers can successfully interact with security engineers and architects to design a comprehensive IA solution.
- **Scientists and researchers.** To focus their efforts on customer requirements not being met by current technology. Thus, the Framework will highlight future IA technology and identify technology gaps for use by both government and commercial research communities.
- **Commercial product and service providers.** To gain insight into the needs of customers. Industry will get an indication of the current and future markets for IA products and services.
- **Standards bodies and consortia.** To provide guidance in developing standards for commercial products. A major emphasis within the customer base focuses on the use of commercial products, which are driven by commercial standards. The IATF highlights gaps in the available standards that will help focus efforts to influence the standards bodies.

1.3 Context

1.3.1 Information Infrastructures Defined

The IATF is based on the concept of an information infrastructure. An information infrastructure comprises communications networks, computers, databases, management, applications, and consumer electronics and can exist at the global, national, or local level. The global information infrastructure is not controlled or owned by a single organization—"ownership" is distributed among corporate, academic, and government entities as well as by individuals. The Internet is an example of a global information infrastructure as is the global telecommunications network. Most organizations that communicate externally rely upon this global system in conducting their operations using a combination of global, virtual networks, dedicated networks, Wide Area Networks (WAN), and customized information systems.

A national information infrastructure is the collection of information infrastructures used by the nation to conduct its business, whether government or commercial. One instance of a national infrastructure is the United States (U.S.) critical infrastructure as defined in Presidential Decision Directive (PDD) 63. Before the growth of multinational companies and the advent of the Internet, one could easily identify a national information infrastructure. In the last few decades however, the lines between the global and national information infrastructures have blurred significantly. Each country will need to decide whether the distinction between the two has merit; if so, criteria will be required to categorize an asset as qualifying as part of a "national" information infrastructure. In the U.S., one possible criterion might be whether assets are subject to U.S. laws, regulations, and policies.

Local information infrastructures are the dedicated assets an organization operates in conducting its business; they consist mainly of commercial information systems, network technologies, and applications. Security measures are applied by the owner or operator of the local information infrastructure—defined either as an organization, or even a business unit within an organization.

1.3.2 Categorizing Information and Information Infrastructures

Within the organization, information processed using these assets are generally grouped into functional categories such as administrative, personnel, or logistics. Some information may be available to the public, some considered private. There are many types of private information; companies have different types of proprietary information, government organizations have many types of classified information, including law enforcement, secret, top secret, and sensitive compartmented information. These divisions of information availability are also called information domains.

To accomplish their various missions and to protect their critical functions, all organizations—both government and private sector—have public and private information they need to safeguard. The mission or business environment determines how, and to what extent, specific information is protected. What is publicly releasable to one organization may be private to another, and vice versa. The Federal Government uses specific categories for some of its private information under the heading of "classified information." In general, the government recognizes four classification levels: unclassified, confidential, secret, and top secret. Within the classification levels, there may be subcategories specific to individual communities. Three of the classification categories—confidential, secret, and top secret—address private information. The fourth level of classification covers both private information (such as sensitive or Privacy Act Information) and public information.

Several types of information could be considered private. One example would be law enforcement information that could potentially damage or impair law enforcement efforts if improperly protected or handled. Proprietary information is much the same for the business community; the information would be harmful to the business if it were released. Information covered under the Privacy Act including personal financial, medical, and other such information

is also considered sensitive. The government handles a variety of classified and sensitive information supporting research, engineering, logistic, administrative, and acquisition functions across the different organizations and agencies.

Most organizations assign more rigorous requirements to protecting their private information than their public information. First access is controlled. For example, within an organization, a human resources or finance person may have complete access to personnel and payroll databases and servers, but may not have access to the most sensitive research and development information. Within the government-classified realm this is accomplished by assigning different classification levels, special compartments, and need-to-know designations. This is illustrated in Figure 1-1.

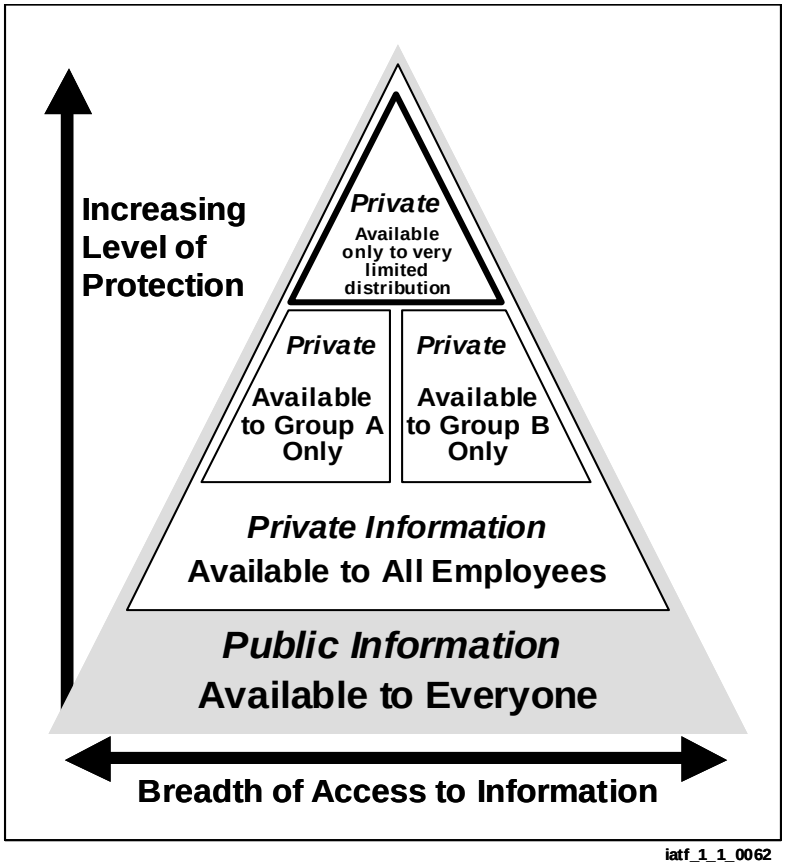

Figure 1-1. Availability and Protection to Information

In addition to access controls, more robust technical security measures are implemented. Organizations acknowledge that the potential loss from exposing private information to the public would be high and therefore the additional cost of protection is warranted. In Figure 1-1, the most stringent security measures would be applied to the information and information infrastructures associated with the top triangle.

The partitioning of information according to access control, need, and levels of protection required yields categories of information. The categories are often called *information domains*. Organizations implement specific mechanisms to enforce the information partitioning and to provide for the deliberate flow of information between information domains.

Protecting information in a collaborative environment presents its own challenges. Organizations sharing information need to agree upon the sensitivity level of the information as well as methods to protect it. Often one organization regards information as more or less sensitive than its partner and officials from both organizations must negotiate a mutually agreeable solution. This occurs between companies sharing proprietary information, between government organizations involved in a joint project, and very often, between countries.

1.3.3 Boundaries and Information Infrastructures

When considering security for information infrastructures, it is important to understand the concept of boundaries. Information assets exist in physical *and* logical locations, and boundaries exist between these locations. An understanding of what is to be protected from external influences can help ensure that adequate protection measures are applied where they will be most effective. However, when analyzing a real-world example, this boundary is not so easily identified. Sometimes the boundary is defined as physical—people, information, and information systems associated with one physical location. But this ignores the reality that, within a single location, many different security policies may be in place, some covering public information and some covering private information.

Other times it is defined as surrounding the information and information systems that are governed by a policy within a single location. This definition, however, does not address the fact that policies cross physical boundaries. Further complicating the matter is that, many times, a single machine or server may house public *and* private information. So, multiple boundaries may exist within a single machine. Figure 1-2 illustrates these complexities associated with defining boundaries. It shows one organization with facilities in two locations, each processing multiple levels of information. In addition, the private network is also connected to the Internet.

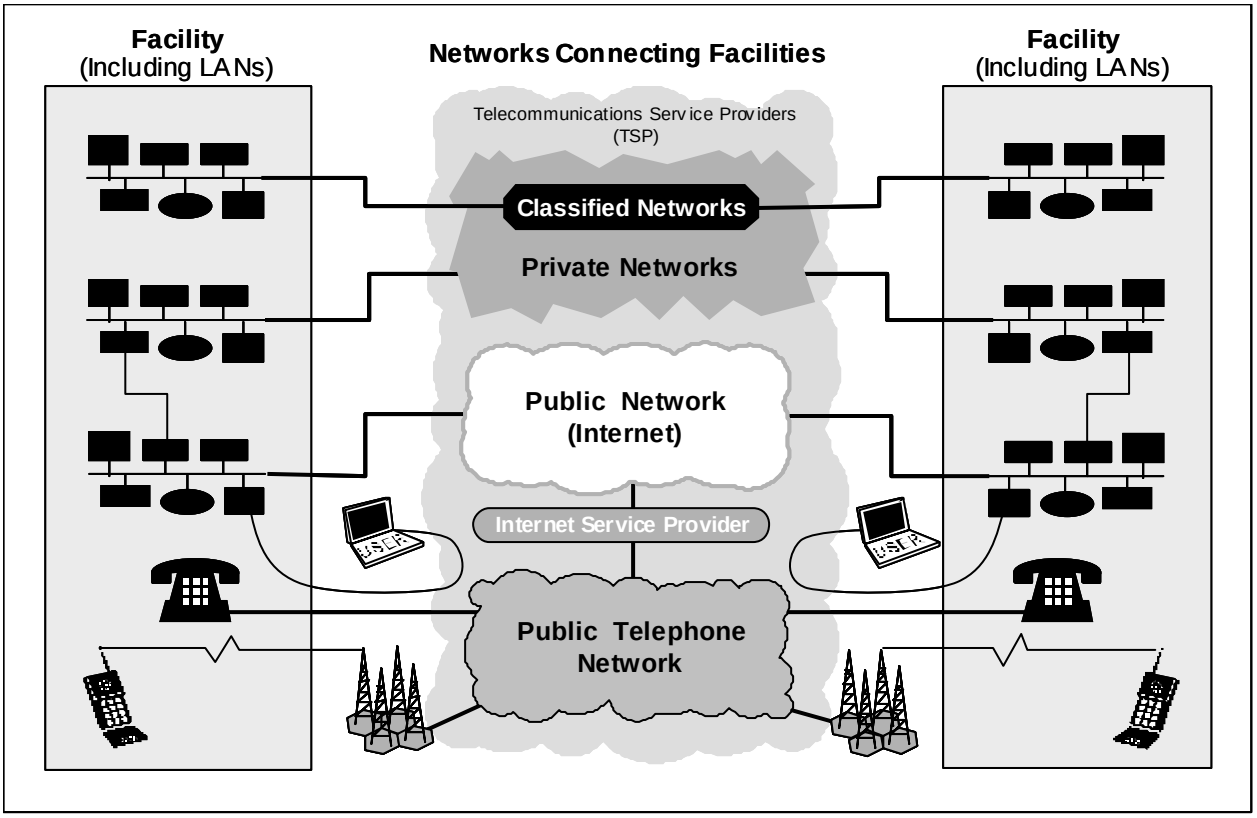

Figure 1-2. Information Infrastructure Elements

In this case, the physical location might be considered a boundary, as might the logical boundaries associated with the different levels of information.

1.3.4 Information Assurance Framework Areas

Given the complexity of information systems, discussion of *how to protect* them is challenging unless a common framework is employed. The IATF document employs a framework that partitions the IA technology aspects of information systems into the following four areas, as shown in Figure 1-3.

1. Local Computing Environments.
2. Enclave Boundaries (around the local computing environments).
3. Networks and Infrastructures.
4. Supporting Infrastructures.

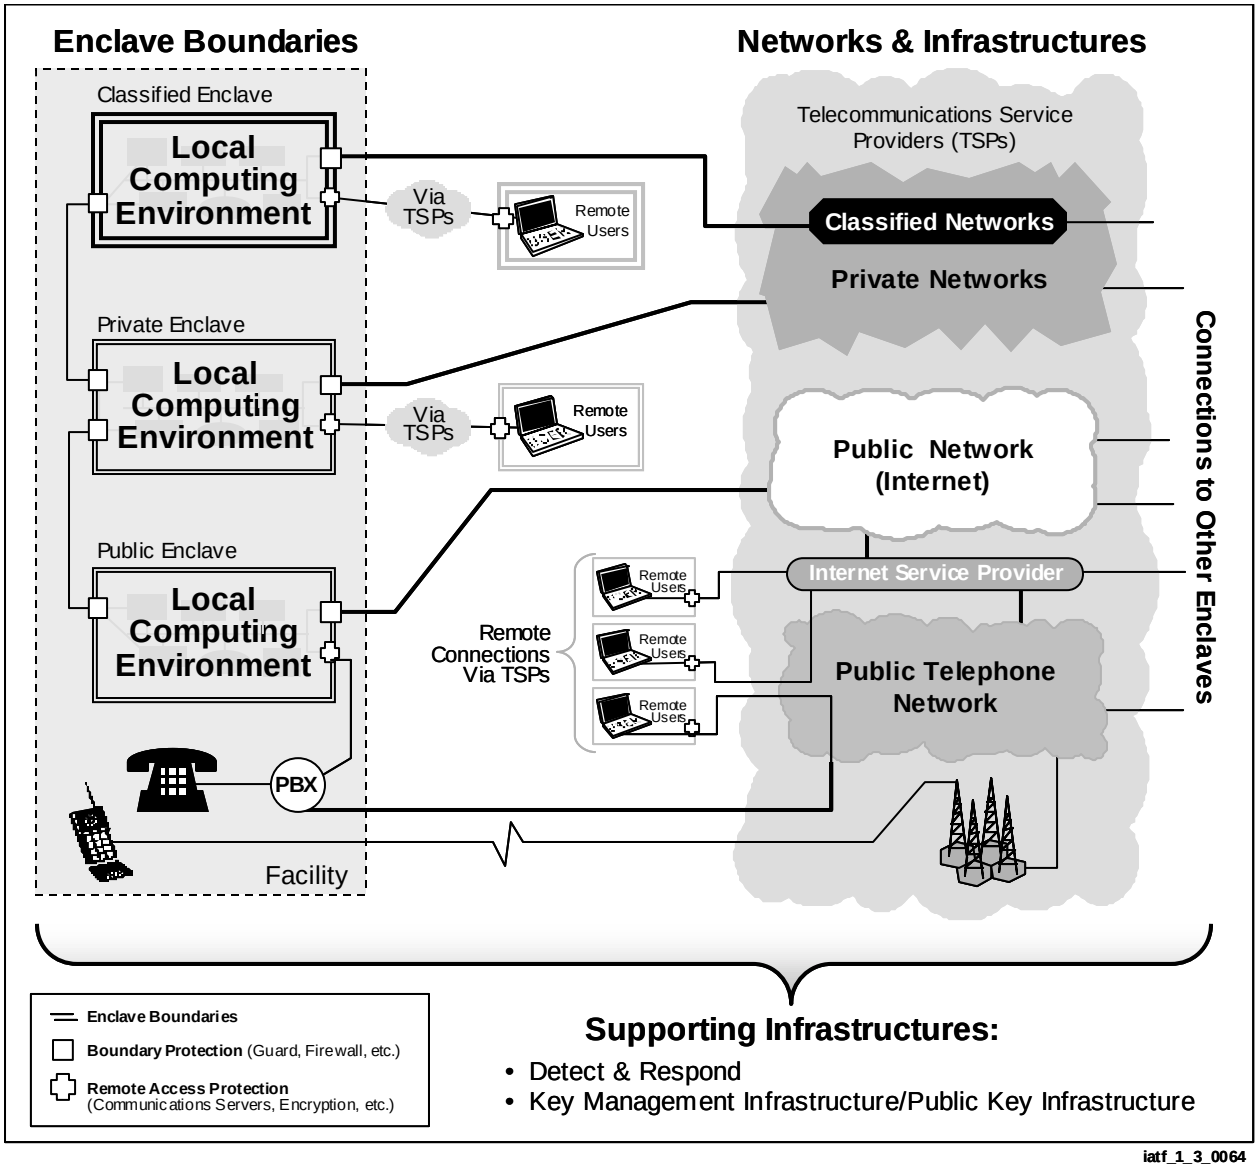

Figure 1-3. IA Technology Framework Access

By partitioning the discussion into these four areas, aspects of IA technology for the information system can be focused upon and more clearly presented. However, these areas are overlapping bins of concern. Effective implementation of IA for a given information system involves the interplay of actions taken throughout the information system—across all four technology framework areas. In the paragraphs that follow, the four framework areas are described further.

Local Computing Environments Framework Area

The local user *computing environment* typically contains servers, clients, and the applications installed on them. Applications include, but are not limited to, those that provide services such as scheduling or time management, printing, word processing, or directories. This environment is represented in Figure 1-4.

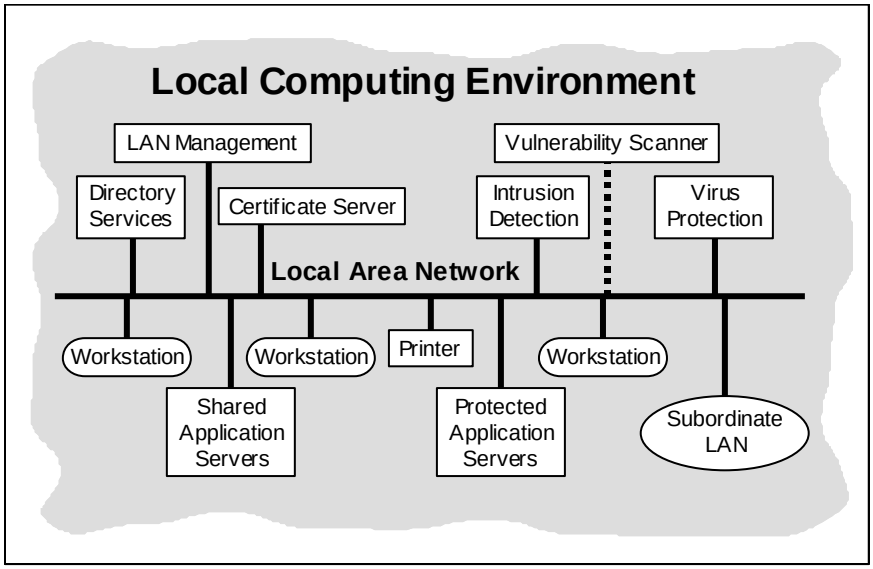

Figure 1-4. Local Computing Environment Area

Looking across the range of computing environments, there are several broad categories of information systems that organizations employ. In both the private sector and the government, one will find large legacy information systems that have been developed over many years and at considerable expense to satisfy unique mission/business needs. These will likely remain in place for some time to come. A large number of organizations have also heavily invested in the use of commercial off-the-shelf (COTS) products or customized versions of COTS information system components and products tailored for their specific use. Organizations using customized products will probably transition to full COTS implementations as the product offerings address their needs more directly.

Most organizations want to use multiple applications to perform their operational mission functions. As a result, users are struggling to integrate the ever-growing range of applications into an effective information processing capability. Each of these applications will place unique requirements on the supporting infrastructure.

Across the range of computing environments, the customer base needs IA solutions in many existing application areas. Security of the computing environment focuses on servers and clients to include the applications installed on them, the operating systems, and host-based monitoring capabilities. Application areas requiring IA solutions include the following:

- Messaging, e.g., electronic mail (e-mail).
- Operating systems.
- Web browser.
- Electronic commerce.
- Wireless access.
- Collaborative computing.
- Database access.

Enclave Boundaries Framework Area

A collection of local computing devices interconnected via Local Area Networks (LAN), governed by a single security policy, regardless of physical location is considered an "enclave." As discussed above, because security policies are unique to the type, or level, of information being processed, a single physical facility may have more than one enclave present. Local and remote elements that access resources within an enclave must satisfy the policy of that enclave. A single enclave may span a number of geographically separate locations with connectivity via commercially purchased point-to-point communications (e.g., T-1, T-3, Integrated Services Digital Network [ISDN]) along with WAN connectivity such as the Internet. These concepts are represented in Figure 1-5.

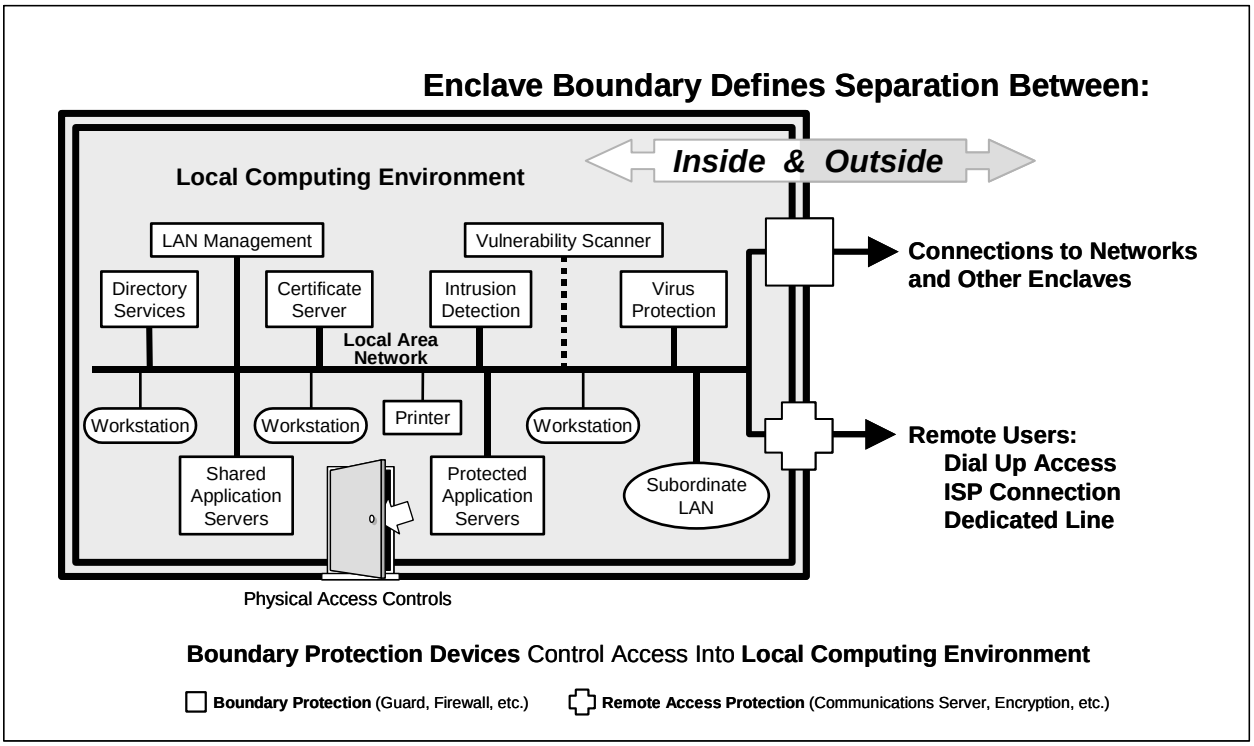

Figure 1-5. Enclave Boundaries Framework Area

The enclave boundary is the point at which information enters or leaves the enclave or organization. Many organizations have extensive connections to networks outside their control. Therefore, a layer of protection is needed to ensure that the information entering does not affect the organization's operation or resources, and that the information leaving is authorized.

Many organizations employ multiple types of external network connections through the enclave boundary. These include:

- Connections to external networks (such as the Internet) to exchange information with another enclave or to access data on a network.
- Three types of connections to remote users—dial-up access via the public telephone network, connection to an Internet Service Provider (ISP) by direct connection (cable modem), or by dial-up access, and dedicated line connectivity through a Telecommunications Service Provider (TSP) (see also Figure 1-3).
- Connections to other local networks operating at different classification levels.

Each connection requires different types of solutions to satisfy both operational and IA concerns. Internets invite access through the boundary, with security only as good as the entire network through which the data is being transported.

Networks and Infrastructures

The network and infrastructure of these networks provide connectivity between enclaves; they contain Operational Area Networks (OAN), Metropolitan Area Networks (MAN), Campus Area Networks (CAN), and LANs, extending coverage from broad communities to local bases. The transport networks contain the information transmission components (e.g., satellites, microwave, other Radio Frequency (RF) spectrum, and fiber) to move information between the network nodes (e.g., routers and switches). As depicted in Figure 1-6, other important components of the network infrastructure are network management, domain name servers, and directory services.

The typical types of transport networks and services used by the government and industry now, and that will be used in the future, can be logically grouped into three areas:

1. Public/commercial networks and network technologies.
2. Dedicated network services.
3. Government-owned and operated.

The public/commercial networks used by the private sector and government include the Internet, the Public Switched Telephone Network (PSTN), and wireless networks. Wireless networks include cellular, satellite, wireless LAN, and paging networks. Access to networks is typically gained through telecommunications service providers. These public networks are wholly owned and operated by these private sector providers.

To obtain dedicated network services, the government has structured a number of network service contracts that procure network services. These include the Federal Wireless Service and

FTS 2000. Public network providers provide access to networks through an arrangement with the government. Private sector organizations obtain telecommunications services in a similar manner, leasing and purchasing dedicated commercial telecommunications services.

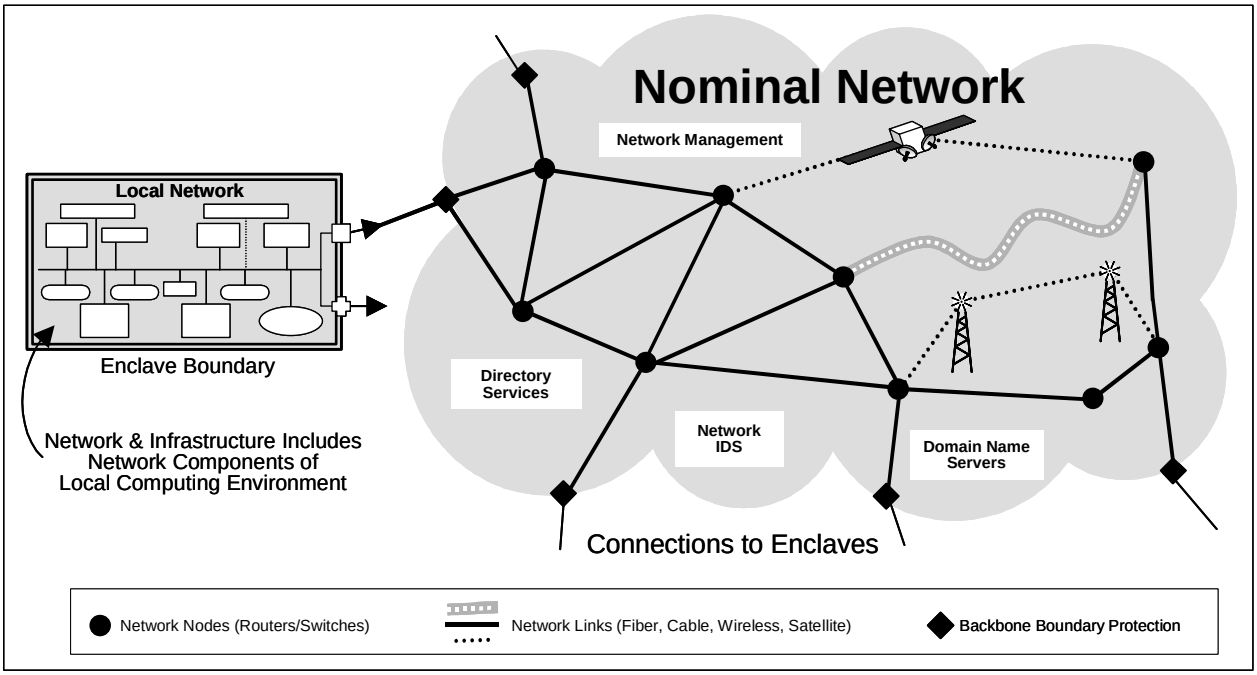

iatf_1_6_0067

Figure 1-6. Network and Infrastructure Framework Structure

Several government organizations own and operate networks. For example, the Department of Energy's Energy Science Network (ESNet), the Federal Aviation Administration's Agency Data Telecommunications Network (ADTN), and the DoD Secret Internet Protocol Router Network (SIPRNET). These networks may begin as private networks, go through leased or public networks, and terminate as private networks. They also include totally owned and operated networks such as MILSTAR. Appendix C provides additional information on this category of networks.

Supporting Infrastructures

Also present in the information technology environment are *supporting infrastructures* that provide the foundation upon which IA mechanisms are used in the network, enclave, and computing environments for securely managing the system and providing security-enabled services. Supporting infrastructures provide security services for networks, end-user workstations, servers for Web, applications, and files, and single-use infrastructure machines (e.g., higher-level Domain Name Server [DNS] services, higher-level directory servers). The two areas addressed in the IATF are key management infrastructure (KMI), which includes Public Key Infrastructures (PKI), and detect and respond infrastructures.

Key Management Infrastructure

A KMI provides a common unified process for the secure creation, distribution, and management of the public key certificates and traditional symmetric keys that enable security services for the network, enclave, and computing environment. These services enable the identities of senders and receivers to be reliably verified, and the information to be protected from unauthorized disclosure and modification. The KMI must support controlled interoperability for users, consistent with established security policies for each user's community.

Detect and Respond

The detect and respond infrastructure enables rapid detection of, and reaction to, intrusions. It also provides a "fusion" capability so one incident can be viewed in relation to others. This allows analysts to identify potential activity patterns or new developments. In most organizations that implement a detect and respond capability, local centers monitor local operations and feed a larger regional or national center. The infrastructure required includes technical solutions such as intrusion detection, and monitoring software; and a cadre of skilled specialists, often referred to as a Computer Emergency Response Team (CERT).

1.3.5 Nature of Cyber Threats

Information systems and networks offer attractive targets. They should be resistant to attack from the full range of threat-agents—from hackers to nation states—and they must limit damage and recover rapidly when attacks do occur.

The IATF considers five classes of attacks:

1. Passive.
2. Active.
3. Close-In.
4. Insider.
5. Distribution.

The key aspects of each class of attack are summarized in Table 1.1.

Table 1-1. Classes of Attack

Attack	Description
Passive	Passive attacks include traffic analysis, monitoring of unprotected communications, decrypting weakly encrypted traffic, and capturing authentication information (e.g., passwords). Passive intercept of network operations can give adversaries indications and warnings of impending actions. Passive attacks can result in the disclosure of information or data files to an attacker without the consent or knowledge of the user. Examples include the disclosure of personal information such as credit card numbers and medical files.
Active	Active attacks include attempts to circumvent or break protection features, introduce malicious code, or steal or modify information. These include attacks mounted against a network backbone, exploitation of information in transit, electronic penetrations into an enclave, or attacks on an authorized remote user when attempting to connect to an enclave. Active attacks can result in the disclosure or dissemination of data files, denial of service, or modification of data.
Close-in	Close-in attack is where an unauthorized individual is in physical close proximity to networks, systems, or facilities for the purpose of modifying, gathering, or denying access to information. Close proximity is achieved through surreptitious entry, open access, or both.
Insider	Insider attacks can be malicious or non-malicious. Malicious insiders have the intent to eavesdrop, steal or damage information, use information in a fraudulent manner, or deny access to other authorized users. Non-malicious attacks typically result from carelessness, lack of knowledge, or intentionally circumventing security for non-malicious reasons such as to “get the job done.”
Distribution	Distribution attacks focus on the malicious modification of hardware or software at the factory or during distribution. These attacks can introduce malicious code into a product such as a back door to gain unauthorized access to information or a system function at a later date.

The relationship of these attack classes to the technology framework areas is shown in Figure 1-7. Subsequent sections of the IATF will provide an overview of the IA strategy for countering or mitigating the effects of these attacks.

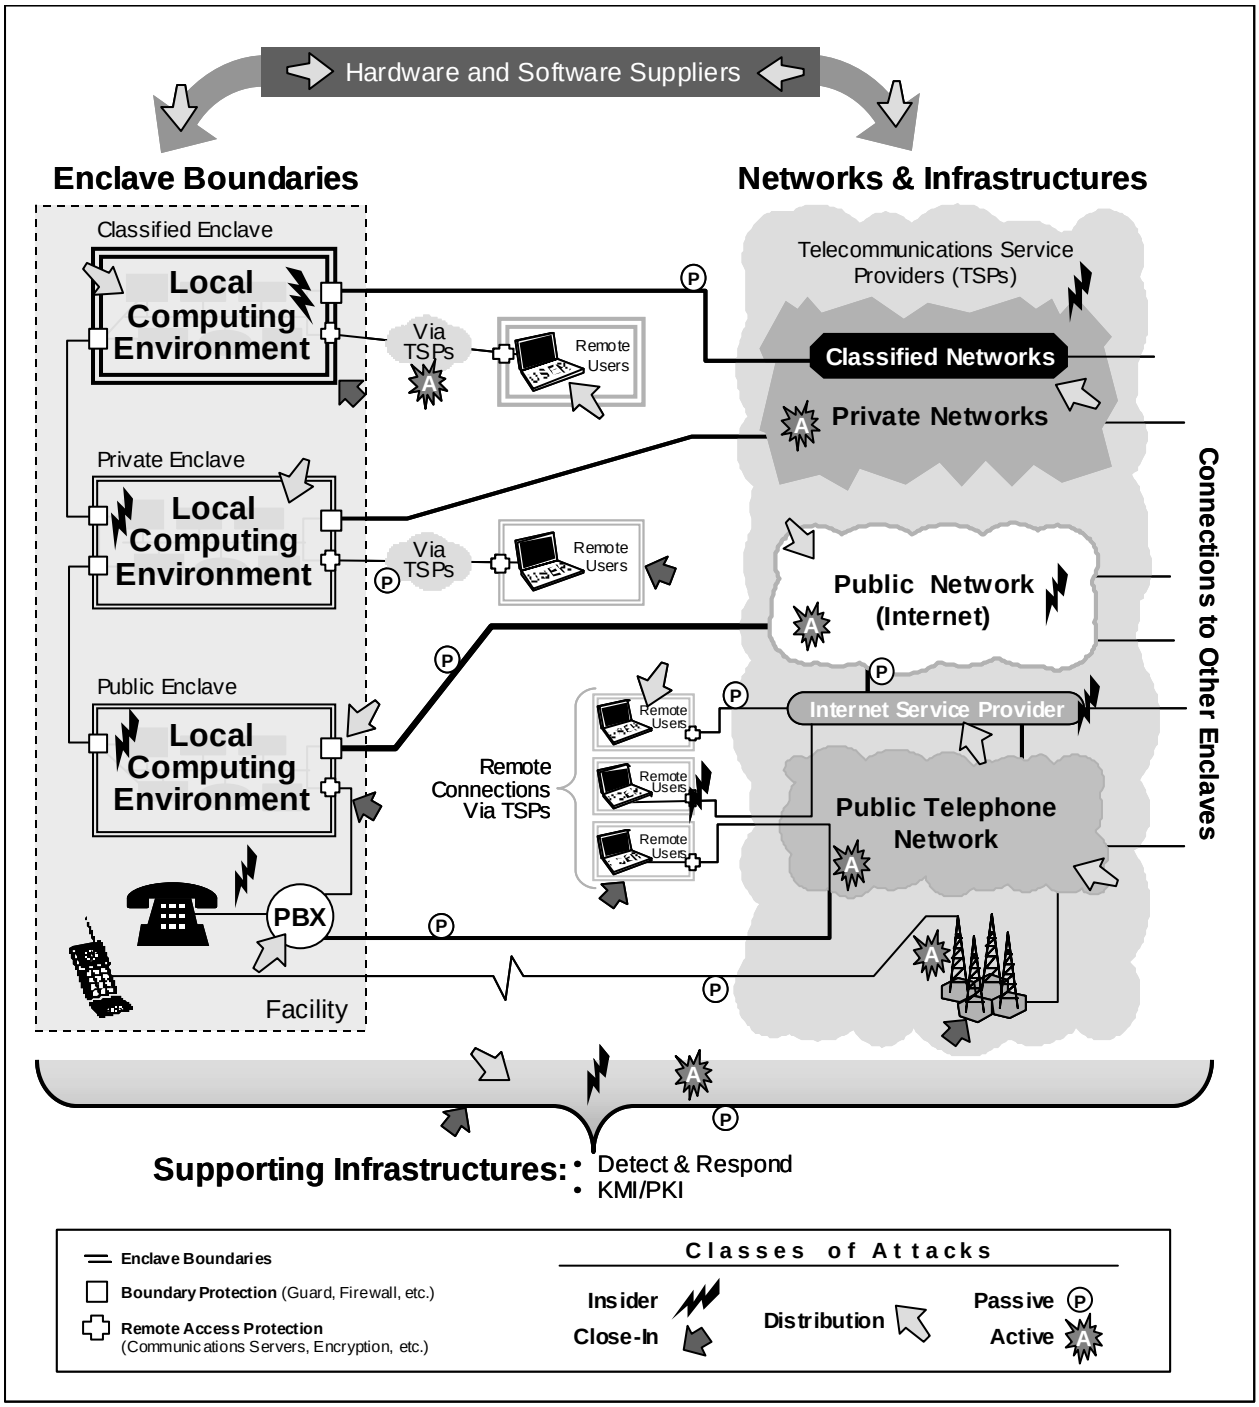

Figure 1-7. Classes of Attacks on the Information Infrastructure

1.4 Defense-in-Depth

The Department of Defense (DoD) has led the way in defining a strategy called *Defense-in-Depth*, to achieve an effective IA posture. The underlying principles of this strategy are applicable to any information system or network, regardless of organization. Essentially, organizations address IA needs with *people* executing *operations* supported by *technology*.

Figure 1-8 illustrates the principal aspects of the Defense-in-Depth strategy—personnel, technology, and operations, outlined as follows.

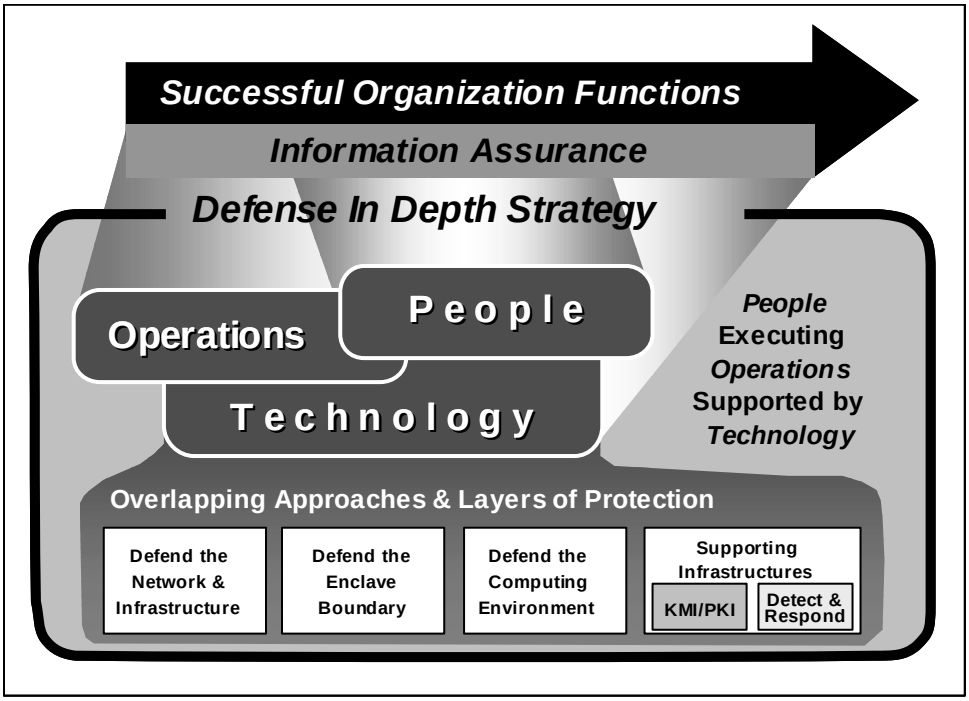

Figure 1-8. Principal Aspects of the Defense-in-Depth

- People
  - Policies and Procedures
  - Training and Awareness
  - Physical security
  - Personnel security
  - System security administration
  - Facilities Countermeasures
- Technology
  - IA Architecture framework areas
  - IA criteria (security, interoperability, and PKI)
  - Acquisition integration of evaluated products
  - System risk assessments
- Operations
  - Security policy
  - Certification and accreditation
  - Readiness assessments
  - Security management
  - Key management
  - Attack sensing and warning response
  - Recovery and reconstitution

Of the three principal aspects of this strategy, the IATF focuses on technology and on providing a framework for providing overlapping layers of protection against cyber threats. By this approach, a successful attack against one layer or type of protection does not result in the compromise of the entire information infrastructure.

Other policies, procedures, and frameworks are focused on addressing the people and operations aspects of a Defense-in-Depth strategy.

1.4.1 Defense-in-Depth and the IATF

Information infrastructures are complicated systems with multiple points of vulnerability. To address this, the IATF has adopted the use of multiple IA technology solutions within the fundamental principle of the Defense-in-Depth strategy, that is, using layers of IA technology solutions to establish an adequate IA posture. Thus, if one protection mechanism is successfully penetrated, others behind it offer additional protection. Adopting a strategy of layered protections does not imply that IA mechanisms are needed at every possible point in the network architecture. By implementing appropriate levels of protection in key areas, an effective set of safeguards can be tailored according to each organization's unique needs. Further, a layered strategy permits application of lower-assurance solutions when appropriate, which may be lower in cost. This approach permits the judicious application of higher-assurance solutions at critical areas, (e.g., network boundaries).

The Defense-in-Depth strategy organizes these requirements into four principle areas of focus:

1. Defend the Network and Infrastructure.
2. Defend the Enclave Boundary.
3. Defend the Computing Environment.
4. Supporting Infrastructures.

These four areas of focus for the Defense-in-Depth strategy parallel the four framework areas discussed in Section 1.3.4.

1.5 IATF Organization

This framework document has been assembled to present the technology aspects associated with the Defense-in-Depth framework areas; each of the four areas is presented in a separate chapter. Also present are chapters that address concerns that cut across the technology areas or address the IA needs of particular environments or technologies.

To focus on the needs of a diverse group of readers, the IATF is organized into four primary parts shown in Figure 1-9:

- Main Body (Chapters 1-4),
- Technical Sections (Chapter 5-10 and Appendices A-E, H-J),
- Executive Summaries (Appendix F),
- Protection Profiles (Appendix G).

The main body of the IATF (Chapters 1 through 4) provides the general IA guidance that information system users, security engineers, security architects, and others can use to gain a better understanding of the IA issues involved in protecting today's highly interconnected information systems and network backbones. The technical sections (Chapters 5 through 10 and Appendices A through E and H through J) provide specific requirements and solutions for each of the Defense-in-Depth areas. The technical sections also offer the government and private research communities a perspective on technology gaps that exist between today's best available protection solutions and the desired IA capabilities.

For users and security engineers looking for more definitive guidance, the Executive Summaries portion of the IATF provides outlines of the threats, requirements, and recommended solutions for a variety of specific protection needs in specific environments. The goal of this collection of Executive Summaries is to offer quick reference guides (each summary is targeted to be fewer than three pages in length) that users and security engineers can peruse to find scenarios similar or identical to their own IA challenges.

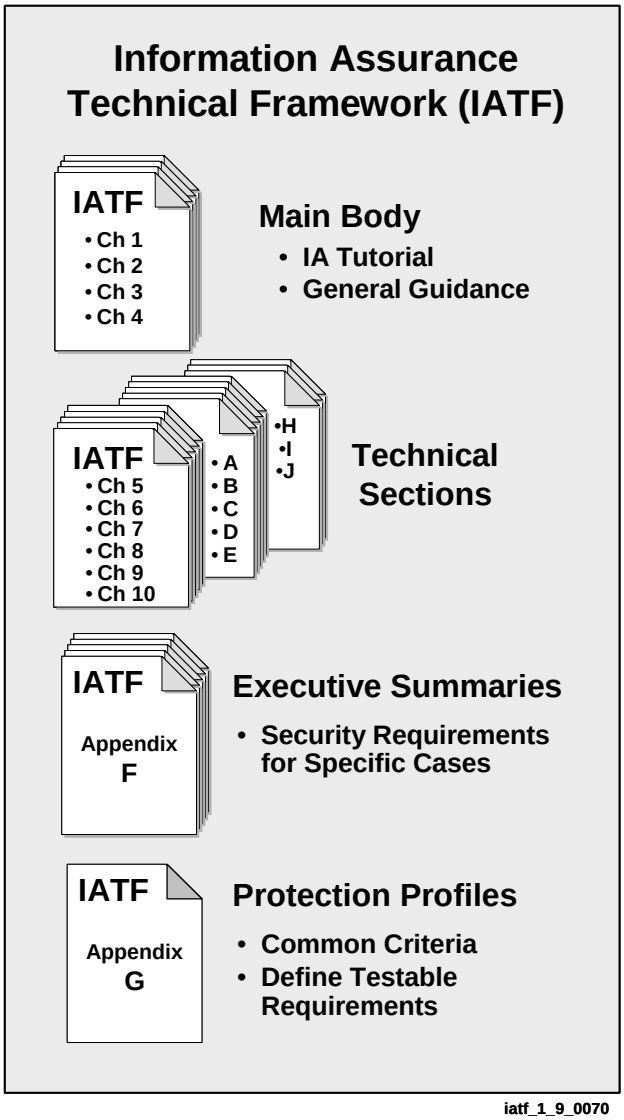

Figure 1-9. Composition of the IATF

Executive Summaries are under development and will be included in a future release of the IATF. For this version of the IATF, an outline illustrating the content of an Executive Summary is provided in Appendix F. In identifying IA solutions, the Executive Summaries will point to the documentation sources (e.g., specifications and protection profiles) containing the set of testable requirements satisfying the user need.

The fourth part of the IATF includes referenced Protection Profiles. Protection Profiles (Appendix G) capture the assurance requirements and functionality for a system or product. They employ the international standard Common Criteria language and structure.

This page intentionally left blank.