SAND REPORT

SAND2002-3558
Unlimited Release
Printed November 2002

Sandia Extended Network:

Design Requirements

Michael D. Gomez and Steven A. Gossage

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94-AL85000.

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865)576-8401
Facsimile: (865)576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.doe.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800)553-6847
Facsimile: (703)605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2002-3558
Unlimited Release
Printed November 2002

Sandia Extended Network
Design Requirements

Michael D. Gomez
Telecommunication Operations Department

Steven A. Gossage
Advanced Networking Integration Department

Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0812

Abstract

This report contains the design requirements for creating a limited-access Sandia Extended Network (SXN), which would be used to collaborate with Nuclear Weapons Complex Labs personnel, university collaborators, industry, and others who may not be allowed accounts on the Sandia Restricted Network (SRN).

Introduction

This document contains the design requirements for creating a limited-access Sandia Extended Network (SXN), which would be used by non-Sandians to collaborate with NWC Labs personnel and others who are not allowed accounts on the Sandia Restricted Network (SRN). Its main purpose is to articulate the requirements upon which the design options and hardware costs for the Sandia eXtended Network (SXN) can be based and in turn presented to 8900 and 9300 Management. The requirements are further addressed in reports outlining its security architecture and in the five-volume set of network architecture reports (*An Architecture for the Sandia Extended Network: Overview; Detailed Description of the Architecture, Design of the Model, and Balanced Protections; Background of the Architecture and its Relevance to Sandia; Terminology and Concepts Relevant to Networks; and Policy-Based Networks and Information Management*.

A Corporate SRN Access Process (CSAP) was chartered in 1999 by Pace VanDevender, Chief Information Officer for Sandia National Laboratories. Its intent was to mitigate the risks associated with having outsiders use the SRN, about whom little was known or could be verified, such as foreign visitors. To mitigate risks associated with these users, creating a non-common, need-to-know environment, such as a limited-access network outside the SRN, was proposed, and the CIO council concurred with it.

The design team expects that by July 31, 2002, the CIO will ratify the Capability Release 1.0 (SXN Design) document. To this end, the SXN design will be fully documented, as well as the design options that were considered, including associated cost estimates for equipment. The team will study these design alternatives and a decision analysis will be conducted to select a final system design. If management concurs with and decides to proceed with the design, detailed staff estimates, deployment details and project planning, then Capability Release 1.0 can be completed, documented and submitted to 8900 and 9300 for final approval under the title “Final System Design Document – Deployment Detail.”

This document contains key concepts, previous design concepts proposed under the auspices of the CSAP, sections on performance drivers and other considerations, the requirements themselves, and then three appendices with background information to support the requirements.

Key Concepts

Access Control – Preventing unauthorized use of a resource (network service, network device) by using some kind of controls to limit access to it.

Audit – Performing an official, independent review and examination of a system, including its records, operational procedures, and system activities to test the adequacy of system control, to ensure compliance with established policies and procedures, and to recommend necessary changes in controls, policies, or procedures.

Audit Trail – A system of record keeping to enable an audit. The audit trail tracks activities sufficiently to enable a reconstruction, review, and examination of the sequence of environments and activities surrounding or leading to each event in the path of a transaction, from its inception to output of final results.

Authentication – The act of verifying the identity of a terminal, workstation, originator, or individual to determine that entity's right to access specific categories of information.

Authorization – The act of granting rights to access a terminal, transaction, program, or process. Authorization is generally used in conjunction with the concept of authentication. Once a user has been authenticated, the user may then be authorized for different types of access within a process, transaction, or program.

Least Privilege – An access principle requiring that each entity be granted the most restrictive set of privileges needed to perform authorized tasks. The application of this principle limits the damage that can result from accident, error, or unauthorized use.

Risk – Risk is a function of the likelihood of a given threat-source exercising a particular potential vulnerability and the resulting impact of that adverse event on the organization.

Risk Assessment – Risk assessment is the first process in the risk management methodology. It is used to determine the extent of the potential threat. The output of the process helps identify appropriate controls for reducing or eliminating risk during the risk mitigation process.

Risk Mitigation – Risk mitigation involves prioritizing, evaluating, and implementing the appropriate risk-reducing controls recommended from the risk assessment process.

Before the CSAP Design Review of March 2001, a proposed high-level design concept was developed that would meet the intent of the CSAP issues. That design concept is shown below:

Proposed Design Concept for Limited Access Network

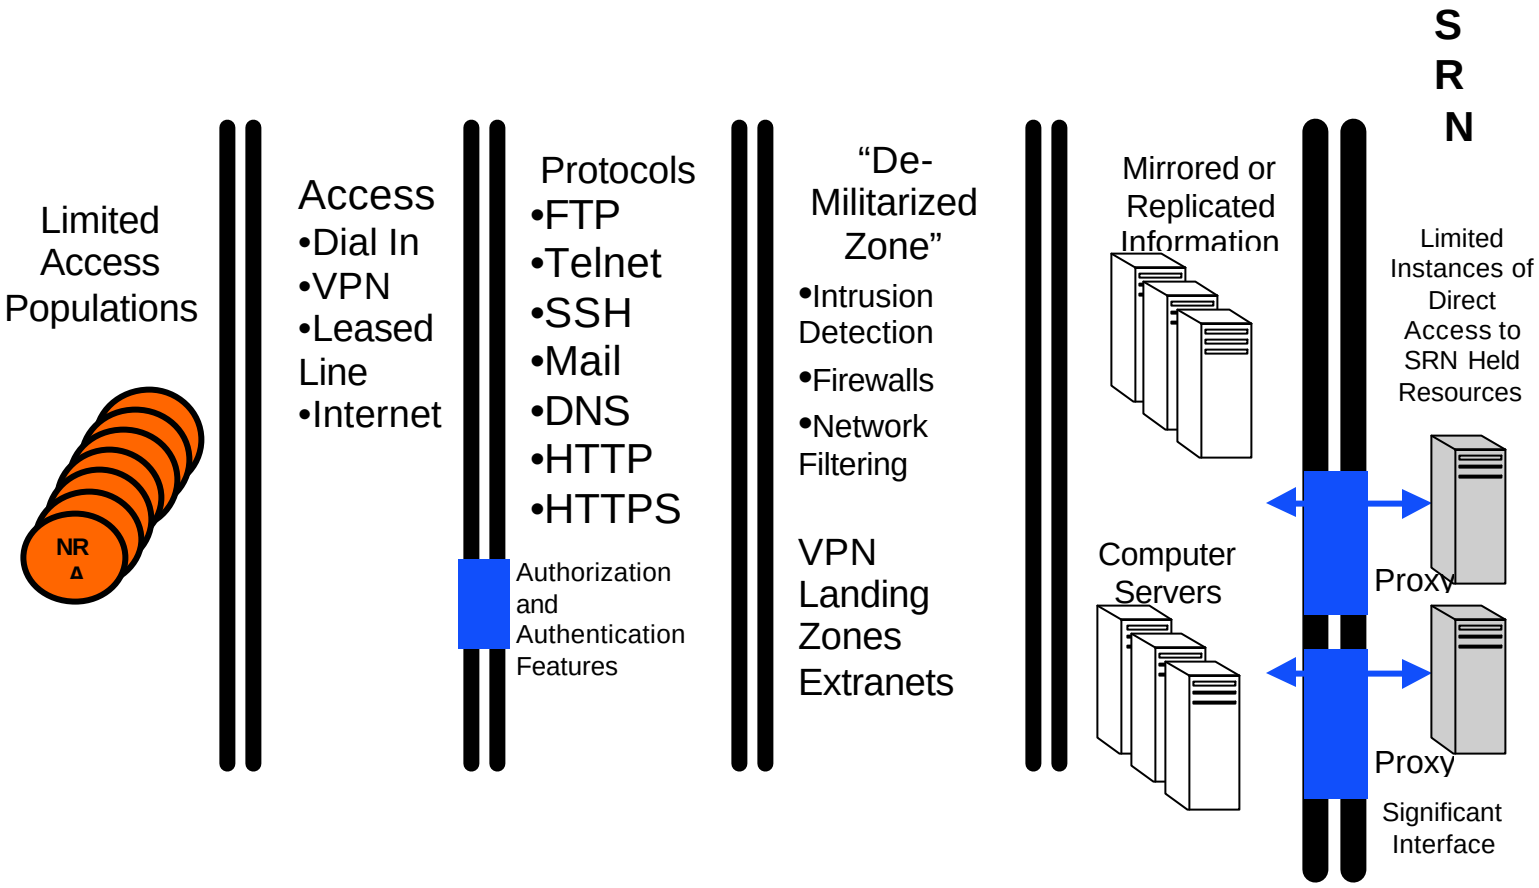

Creation of a known high-level, end-to-end future state that is

- known
- extensible
- verifiable
- compliant

Previous Concepts Considered

In March 2001, Deputy CIO Doug Weaver proposed the five options shown below for the CSAP for the CIO and the IIS Directors to consider. The three options approved for further study and carried forward from that discussion are shown in boldface. In May 2001 a group of IIS Architects studied these options for solving issues related to the CSAP.

1. **Continue to implement dual-intent SRN**
   - Internal open system, common need-to-know environment
   - External non-common-need-to-know environment
2. **Separate SRN common need-to-know and non-common-need-to-know network domains and services.**
3. Implement the SRN as only a non-common-need-to-know environment
4. **Implement the SRN as only a common-need-to-know environment; move NCNTK to the SON.**
5. No action

The architects overwhelmingly selected what has come to be known as "Option 2" using a matrix developed by Mike Gomez, 9334, which was a modification of the matrix previously used to study the issue of removable media. The matrix comprises 31 elements grouped into six broad categories. A representation of the categories studied and summary of the analysis results are shown below:

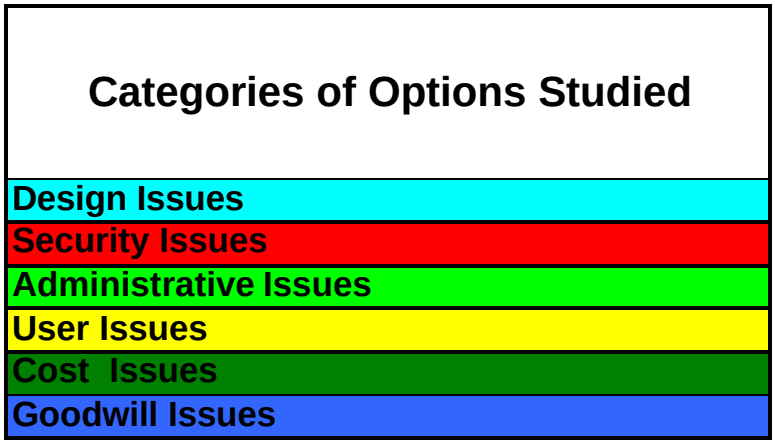

Title of Solution	Description of Solution		Option Average Scores
Dual Intent SRN	Internal open system with Common Need-to-Know environment and external Non-Common Need-to-Know environment		
	Respondent	**Brown**	**3.10**
		Cuyler	**3.03**
		Gay	**2.03**
		Gossage	**1.00**
	Standard Deviation		**1.25**
	Average Score		**2.32**
Separate SRN CNTK and NCNTK Environments	Separate SRN into Common Need-to-Know and Non-Common-Need-to-Know domains and services		
	Respondent	**Brown**	**3.90**
		Cuyler	**3.94**
		Gay	**3.00**
		Gossage	**3.13**
	Standard Deviation		**0.74**
	Average Score		**3.52**
CNTK on SRN and NCNTK on SON	Implement the SRN as only a common need-to-know environment and move Non-Common Need-to-Know to SON		
	Respondent	**Brown**	**3.16**
		Cuyler	**2.52**
		Gay	**1.87**
		Gossage	**1.00**
	Standard Deviation		**1.14**
	Average Score		**2.15**

As can be seen, Option 2, “Separate SRN CNTK and NCNTK Environments,” was selected by a wide margin. A graphical representation for that option is shown below:

"Option 2"

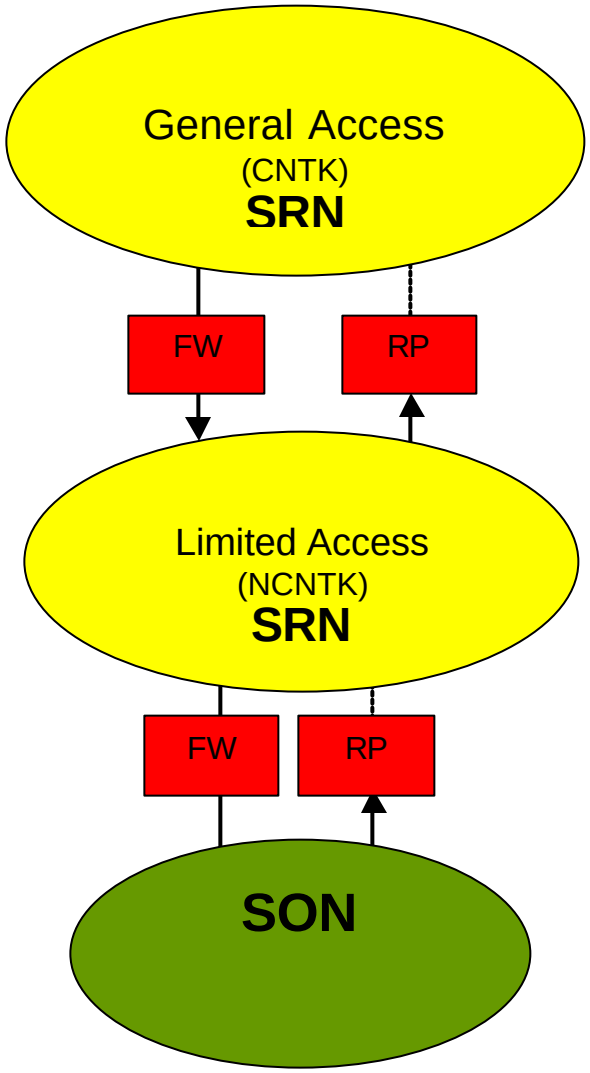

Separate SRN GA and LA Environments

Note:

This diagram shows the logical separation of the SRN CNTK and NCNTK environments, not physical connections between them (i.e., access between the SON and General-Access SRN will not be through the Limited-Access SRN, now known as SXN.)

Definitions

General-Access SRN: The capability of a user to access all SRN resources that are not protected by their own access controls (i.e., host-controlled or internal-network-controlled access). Access will take place through authorized SRN access channels.

Limited-Access SRN: The capability of a user to access only a limited set of SRN resources specifically authorized for the user or to groups to which the user belongs. Access will take place through authorized SRN access channels and will be controlled by SNL-approved authentication and authorization mechanisms.

Separate SRN GA & LA network domains & services

Performance Drivers and Capacity Requirements

A major performance driver for the restricted networks is security and need-to-know as expressed by the requirements for appropriate access and auditable confidence, corporate rules, and DOE regulations. Sandia must also protect UCI information, and in addition collaborate with many external agencies, universities and side-by-side contractors. It is also possible for the populations involved in these collaborations to grow and shrink rapidly. This limited-access environment, envisioned as the SXN, must be able to handle roughly 500 limited-access individuals as a nominal population but should be designed to handle up to 5,000. In addition, the environment must be robust enough to maintain a viable relationship with the SRN, which connects over 25,000 ports in Albuquerque, NM and Livermore, CA.

Significant Considerations

Compliance with information-protection requirements for UCI, assuring appropriate access and auditable confidence, risk identification, risk mitigation, balanced protection and consistent guidance were the key considerations in creating a description of a network architecture as part of the system of protection elements. The network architecture assumes that

- Local on-site desktop SXN users may be able to “sniff” the traffic on their local network and system mitigations are in place

- Local on-site users who have physical access to the general access SRN infrastructure (computers, communications drops, etc.) have been instructed in their responsibilities and the consequences of violations. Accidental or willful actions outside proper use of the SXN are detected by a variety of network scans, switching infrastructure tests, etc., designed to detect inappropriate actions.

- Local on-site SXN users are executing according to the rules, and excursions from the rules are accidental. *The on-site SXN users are not considered to be trusted insider threats (otherwise physical access and infrastructure would have to be controlled in a much more ironclad manner).*

- User desktops, local or remote, are separated from the information servers through distinct, auditable boundaries that are configured to exclusively allow only limited, specific functions

- SXN servers and client machines fall under specific, strict configuration management according to CSAP requirements.

❑ The general-access SRN user desktops have independent boundaries, with SXN servers and user desktops that recognize and enforce appropriate trust relationships.

Assumptions

- General-access users have access to the same information that they currently have, including information residing on limited-access resources
- Need-to-know governs all access by limited-access users
- The sponsoring manager approves non-employee access

Requirements

High-level requirement levied by CSAP Project

The requirement is to deliver a new demarcation zone (DMZ) and a new environment called the Sandia eXtended Network consisting of SXN desktop and SXN server network partitions that are compliant with the requirements of CSAP, the Cyber Security Architecture and the SXN Network Architecture.

Requirements imposed by the Network Architecture
A limited-access network (such as the SXN) must:

- ❑ Provide a consistent security posture for internal and remote access that meets the requirements for transmitting unclassified controlled information (UCI).
- ❑ Function as a part of a balanced, interlocking, consistent, layered protection system that encompasses the information path and control from the user to the host system
- ❑ Enable appropriate collaboration in a secure fashion with auditable confidence
- ❑ Effectively enable need-to-know protections

Additional Requirements:
- Limited-access users reside in the SRN Limited Access Network (SXN), SON, other sites, and the Internet
- SXN users may not be given exceptions or general SRN Access
- Limited-access resources reside in limited-access network unless “proxied” from another network
- The resource owner determines the environment a resource resides in
- Boundaries between SRN general access and limited access are discrete, independent, and; when practical, physical
- Strong interface protections (firewalls, etc.) exist between general-access and limited-access environments
- The resource owner is responsible for applying access controls to the host prescribed by the IIS
- The IIS offers resource owners a service to register limited-access resources
- The SXN requires distinct separation of limited access users and resources

- Absolutely no exceptions can be made to grant limited-access users general access, otherwise revert to current state.
- Adopting the SXN may require duplication of resources and network access elements (increases cost)

Appendix A

Limited-Access (SXN)/General-Access Functionality Today: What is being done

1) Limited-access persons/computers are grouped coarsely
2) Global (coarse-grain) access-control lists (ACLs) are applied to the limited-access person/computer group
3) The limited-access person/computer will be successfully and reliably connected to the target information or development server, potentially with full, general privileges
4) The limited-access person may be located onsite or offsite
5) Authorization is manually performed; case by case

Limited Access/General Access: Liabilities of current state (July 2002)

1) Coarse grouping means that limited-access persons/computers may have access to each other (within their "home" VLAN) although they have no requirements regulating how they do so. Basically, the grouping is coarse enough that all of them are in a common-need-to-know group as a starting point.

2) Coarse intermediate filters (ACLs) cannot accept user identity as a control mechanism. ACLs do not scale well and are difficult to manage centrally. The result is low auditable confidence. ACL logging is weak.

3) Information servers are not consistently configured so that a limited-access person/computer could use a function on the server such as Telnet to try to connect to other hosts in the general-access network. The result is low auditable confidence.

4) The concept of limited privilege is not applied. Limited-access users may obtain more functionality than is required to perform the task at hand.

5) The current implementation is scalable to a few hundred users with difficulty.

Limited Access/General Access: What needs to be done

These are the gaps between the current state and the intermediate and future target states for limited-access and general-access restricted networks.

1) Uniform authentication of limited-access persons

2) Audit of limited-access utilization

3) Documentation of current state of limited-access network and server access, network and server configurations

4) Consistent configuration of information servers provided to limited-access persons/computers

5) Due to funding constraints and technical limitations, there is no log of ACL execution (who, what, when, ports, addresses, protocols, etc.) to verify that the provided access is being appropriately used. There is no central documented view of all of the ACLs used to provide limited access. As a result, auditable confidence is low.

6) Attachment methods for limited access are dependent on people and their current responsibilities and not on repeatable processes. An implementation/operational record of who was connected, what was connected, circuit and server configurations, reason for connection, is not being kept. People and identity policies are needed as well as a continuous infusion of relevant technologies. Policy development has a back seat with respect to the research for and development of technology.

Appendix B

Network Design Principles (from Network Architecture)

- Integrate (design) security into the system from the start

- Note the contrast between “designing out” the inherent connectivity features, services, and capabilities of networks, servers, and information systems and “designing in” security and appropriate access.

- The SXN should be designed as a least privilege system that offers the least amount of privilege necessary to perform a given task. The concept of least privilege is independent of the specific protections that may be available. It is founded on the proper allocation of privileges, permissions, authorizations, etc., given to individuals, groups, and other entities.

- The principle of least privilege should be seen as a sharp contrast to the broad-based access and privilege afforded the current general-access SRN, where the

security-protection model is more nearly described as a defense-in-depth system where multiple, overlapping security mechanisms are in play.

A high-level, service-oriented policy must be defined for the limited-access network (SXN) that expresses access control requirements from a customer perspective (in plain English that the user can understand). A low-level security policy must be defined for the SXN that expresses access-control requirements from an implementation perspective. The two policies must be consistent.

- Implementations will change and adapt as network technologies change and desired services evolve. Service and security policies must be independent of the implementations selected. An example of a service policy could be a list of permitted incoming services or a list of permitted outgoing services. The design teams must define what goes on the list. Some prickly services include FTP, Telnet, SSH, email, DNS, etc. Many more policies are a necessary foundation. Note that policy definitions do not constitute business rules, but the policy definitions should be consistent with business rules.

- The design must make the inter-network boundaries clear. Each environment must define where controls will be applied, what controls are applied, to whom the controls are applied. The design challenge is assuring the security of a collection of interconnected networks (inter-network access.)

- The SXN must provide confidentiality in the sense that information will only be provided to authorized individuals. (As an example, pings to data repositories from SXN local and remote desktops may be blocked.) Access control may be performed by a target information system or access control may be performed on behalf of a target information system. In the latter case, the access-control system would be independent of the target information system. It is valuable to be consistent and uniform in the application of access controls. When the location of the selected access control mechanism shifts with the type of access desired, the complexity of the system grows, possibly at the expense of appropriate access.

- The SXN must define the services it will offer. Furthermore a process is needed to remove old services that are no longer needed and to add new services.

- The SXN must define the permitted forms of access.

- Multiple subnets must be designed with distinct functional responsibilities. An example would be a subnet where all access control is performed.

- The inter-networks should be designed with both vertical and horizontal service scaling in mind. The system design needs to target a customer population size and be able to adapt to increases.

- Services should be decomposed into multiple smaller services to facilitate operation, availability, and security.

- The network should contain access networks, DMZ(s), services networks, and information storage networks with appropriate access controls in between.

- The SXN should categorize the desirable features of various generic barriers; for example, proxies, application gateways, connection gateways, tunnels, host-based protections, bastion hosts, etc.

- The SXN should be built according to a network model that differentiates it from the current SRN.

Appendix C

Issues Related To Handling Sensitive Information

The requirements for proper handling of unclassified controlled information (UCI) and unclassified controlled nuclear information (UCNI) are well described in the *Computer Security Desk Reference* [CSDR], in DOE Order M 471.1-1, and in documents on the Department 3132 (Classification & Information Security) Home Page. Numerous Corporate Process Requirements (CPRs) including CPR400.3.5 "Unclassified Foreign Visits and Assignments", and CPR400.2.13.1 "Access to the Sandia Restricted Network (SRN)" reinforce the roles and requirements for the CIO, users, and the IIS as responsible stewards for network access. It is clear from these documents that the limited-access network must support and enable proper need-to-know. It is also clear that one of the design challenges for the network will be the facilitation of numerous *mutually exclusive need-to-know interactions* with at least twelve individual classifications of UCI information.

Least Privilege Principle. An individual should be granted enough privilege to accomplish assigned tasks, but no more. This principle should be applied in direct proportion and with increased rigor as the potential for damage to a system rises.

Least Access Principle. All data owners must grant a particular kind of access for it to become available. The application of this principle limits the damage that can result from accident, error, or unauthorized use.

A limited-access network architecture should be designed according to a recognizable structure and guided by a set of principles. Key design principles for a limited-access network are (a) least privilege, and (b) least access. In addition, the architecture design must be based on clear security policies. The differences between the SRN and the SCN are described the Computer Security Desk Reference, sections 7.11 Unclassified LANs and 7.12 Classified LANs. Excerpts from those sections follow.

Requirements for unclassified LANs (such as the SXN) are

- Owners of a LAN must determine the sensitivity of the information on the LAN and decide whether they need to be in the SRN or SON. If the LAN contains sensitive information and is to reside in the EOE (External Open Environment), the LAN must be protected with a firewall. Once this is done, the LAN becomes an island of protection, which places it under the auspices of the Internal Restricted Environment (IRE).

- All unclassified LANs are covered under the Unclassified Master Plan and Center Information Plans, which describe the sensitivity of the information residing on the LAN and the security measures in place to protect the information. No system connected to a LAN may be simultaneously connected to a system in any other environment.

- If all the users of a LAN do not have a common need-to-know for all the information on the LAN, the LAN must implement passwords, discretionary access controls, and logging of failed logons and attempted invalid accesses.

- If a LAN connects to another LAN whose users do not all have need-to-know for its information, it must implement sufficient mechanisms to prevent unauthorized users from accessing the information.

- No system, while attached to an unclassified LAN, may process, print, or store classified information.

Selected Bibliography

(represents a selection of relevant documents consulted)

"Access to the Sandia Restricted Network (SRN)," sponsor: Pace VanDevender, 9400, Issued: February 1, 1999, Revised: November 1, 2000. (CPR400.2.13.1)

C.D. Brown, M.J. Ernest, L.L. Fine, R.G. Hawkins, R.M. Jansma, L. G. Pierson, W.H. Rahe, L. Stans, "Network Security Design Project Requirements Document," August 25, 1989, Revised December 1, 1989

Cisco SAFE: A Security Blueprint for Enterprise Networks White Paper," Cisco Systems, 2000, 1-65.

Computer Security Desk Reference, July 2001, Computer Security Control Number SNL-DR

David F. Beck, *Developing Design Principles for Information Technology Security,* SAND2002-0453, February 2002.

Identification and Protection of Unclassified Controlled Nuclear Information Manual, DOE M 471.1-1,U.S. Department of Energy, 6/30/00.

"Improving Computer Security through Network Design," Danny Smith, Technical Director AUSCERT, 1997, http://www.auscert.org.au/Information/Auscert_info/Papers/Security_Domains.htm

Information Security Policy for Administrative Information, http://www.aits.uillinois.edu/security/securestandards.html

"Unclassified Controlled Information," http://www-irn.sandia.gov/ADC/1handbook/class/cover010.html

Distribution

MS 0630	M.J. Murphy, 9600
MS 0803	H.L. Pitts, 9601
MS 0622	D.S. Rarick, 9623 (6)
MS 0629	P.D. Merillat, 9500
MS 0801	A.L Hale, 9300
MS 0801	M.J. Sjulin, 9330
MS 0801	W.F. Mason, 9320
MS 0806	C.R. Jones, 9322 (6)
MS 0806	C.D. Brown, 9322
MS 0363	C.A. Morgan, 9323
MS 0822	C. Pavlakos, 9326
MS 0813	R.M. Cahoon, 9327 (6)
MS 0805	W.D. Schwartz, 9329 (6)
MS 0805	D.J. Bragg Jr., 9329
MS 0805	M.W. Gutscher, 9329
MS 0812	M.J. Benson, 9334 (6)
MS 0812	M.D. Gomez, 9334 (10)
MS 0806	Len Stans, 9336 (6)
MS 9003	K.E. Washington, 8900
MS 9011	T.J. Toole, 8910
MS 9012	R.D. Gay, 8930
MS 9018	Central Technical Files, 8945-1
MS 0899	Technical Library, 9619 (2)
MS 0612	Review/Approval Desk, 9612 for DOE/OSTI