

AFRL-IF-RS-TR-2002-193
Final Technical Report
August 2002



CONCURRENT INFORMATION ASSURANCE ARCHITECTURE

WetStone Technologies, Incorporated


APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK**

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2002-193 has been reviewed and is approved for publication.

APPROVED:



ROBERT J. VAETH
Project Engineer

FOR THE DIRECTOR:



WARREN H. DEBANY, Technical Advisor
Information Grid Division
Information Directorate

REPORT DOCUMENTATION PAGE			<i>Form Approved</i> <i>OMB No. 074-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE AUGUST 2002	3. REPORT TYPE AND DATES COVERED Final Apr 01 – Feb 02	
4. TITLE AND SUBTITLE CONCURRENT INFORMATION ASSURANCE ARCHITECTURE			5. FUNDING NUMBERS C - F30602-01-C-0061 PE - 62702F PR - OIAG TA - 32 WU - P2	
6. AUTHOR(S) Hampton Powers and Milica Barjaktarovic				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) WetStone Technologies, Incorporated 273 Ringwood Road Freeville New York 13068			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/IFGB 525 Brooks Road Rome New York 13441-4505			10. SPONSORING / MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2002-193	
11. SUPPLEMENTARY NOTES AFRL Project Engineer: Robert J. Vaeth/IFGB/(315) 330-2182				
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.				12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 Words) Currently, many tasks in the information security field are accomplished in a sequential manner, often after the fact, which limits the urgency of time response and usefulness of the tools and approaches currently available. The next step toward more secure networking is taking a Concurrent Information Assurance (C-IA) approach, which executes security-critical functionality concurrently on several different levels. The C-IA Architecture (C-IAA) postulates concurrent information assurance (IA) by providing configurable, coordinated, automated situation analysis, decision assistance, and response. The C-IAA's objective is to create the underpinnings for an architecture that executes security-critical functionality in an automated fashion, distributively, concurrently, and separately from other applications. C-IAA systems exploit the severability of concurrent processing into separate execution environments to achieve a high confidence and minimal impact on information, IA components, and the organizations dependant on that information.				
14. SUBJECT TERMS Information Assurance, Concurrent Processing, Architecture, High Confidence, Security			15. NUMBER OF PAGES 104	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Contents

1	<i>Introduction.....</i>	<i>1</i>
1.1	Document	1
1.1.1	Purpose.....	1
1.1.2	Overview	1
1.2	Project	1
1.2.1	Background	1
1.2.2	Objectives	2
2	<i>Accomplishments</i>	<i>2</i>
2.1	Tasks and Products	2
2.1.1	C-IA Architecture and Technologies	2
2.1.2	OversightNet Secure Kernel	3
2.1.3	OverNet.....	3
2.1.4	OuterNet.....	3
2.2	Findings.....	3
2.3	Recommendations	4
3	<i>Research Products.....</i>	<i>5</i>
3.1	DARPA Projects Relevant to C-IAA	5
3.1.1	Program: TIDES	7
3.1.2	Program: RKF	7
3.1.3	Program: Information Management	8
3.1.4	Program: Software Enabled Control	9
3.1.5	Program: MICA	9
3.1.6	Program: JFACC.....	9
3.1.7	Program: DASADA	10
3.1.8	Program: ANTS	10
3.1.9	Program: Active Networks.....	10
3.1.10	Program: CHATS.....	11
3.1.11	Program: Dynamic Coalitions	12
3.1.12	Program: Fault-Tolerant Networks	13
3.1.13	Program: NMS	13
3.1.14	Program: Sensor Information Technology	14
3.1.15	Program: EELD.....	14
3.1.16	Program: Command Post of the Future.....	14
3.1.17	Program: Cyber Panel	15
3.1.18	Program: Dynamic Coalitions	15
3.1.19	Program: Fault-Tolerant Networks	15
3.2	Secure Communications: State of the Art	16
3.2.1	Virtual Private Networks	16
3.2.2	Secure Transport Protocols	17
3.2.3	Authentication Technologies	21
3.2.4	Interoperability through Standardization	26
3.3	Intrusion Detection, Analysis, and Response: State of the Art.....	31
3.3.1	Intrusion Detection.....	31
3.3.2	Knowledge Extraction.....	33
3.3.3	Situation and Risk Assessment	37
3.3.4	Expert and Knowledge-based Systems	40

3.4 C-IA Architecture	48
3.4.1 Concept	48
3.4.2 C-IA Requirements	50
3.4.3 Implementation Issues.....	51
3.4.4 Recommended Research	51
3.5 C-IAA Processing Node	54
3.5.1 Concept	54
3.5.2 C-IA Requirements	54
3.5.3 Current State of the Art.....	55
3.5.4 Implementation Issues.....	58
3.5.5 Recommended Research	58
3.6 C-IAA OversightNet	60
3.6.1 Concept	60
3.6.2 C-IA Requirements	60
3.6.3 Current State of the Art.....	60
3.6.4 Implementation Issues.....	61
3.6.5 Recommended Research	62
3.7 C-IAA OverNet.....	65
3.7.1 Concept	65
3.7.2 C-IA Requirements	65
3.7.3 Current State of the Art.....	66
3.7.4 Implementation Issues.....	66
3.7.5 Recommended Research	67
3.8 C-IAA OuterNet	69
3.8.1 Concept	69
3.8.2 C-IA Requirements	69
3.8.3 Current State of the Art.....	69
3.8.4 Implementation Issues.....	70
3.8.5 Recommended Research	78
3.9 Decision-making in C-IAA	81
3.9.1 Stages and Tiers	81
3.9.2 OversightNet Decision-making Example.....	83
3.9.3 Decision Engine Stage I.....	83
3.9.4 Decision Engine Stage II.....	84
4 References	87

Figures

FIGURE 1 — <i>DECISION TREE EXAMPLE</i>	42
FIGURE 2 — <i>DIFFERENT PATHWAYS TO THE SOLUTION, AS SUGGESTED BY [KASA98], P. 66.</i>	47
FIGURE 3 — <i>C-IAA ABSTRACT BLOCK DIAGRAM</i>	48
FIGURE 4 — <i>THE ISO LAYERED PROTOCOL MODEL</i>	73
FIGURE 5 — <i>PROTOCOL ENCAPSULATION USING TCP/IP AND TELNET AS EXAMPLES</i>	73

Tables

TABLE 1 — <i>DARPA PROJECTS SEEKING COMPATIBLE FUNCTIONALITY</i>	6
TABLE 2 — <i>SECURITY-RELATED IETF WORKING GROUPS</i>	28
TABLE 3 — <i>WIDELY USED KDD SOFTWARE</i>	37
TABLE 4 — <i>APPROACHES AT A GLANCE</i>	46
TABLE 5 — <i>MAIN EXCHANGE POINTS ON THE INTERNET</i>	71
TABLE 6 — <i>SAMPLE INTERNET CONNECTION COSTS</i>	75
TABLE 7 — <i>REPRESENTATIVE CONNECTION COSTS</i>	75
TABLE 8 — <i>REPRESENTATIVE CONNECTION COSTS</i>	76
TABLE 9 — <i>ONE-TIME SETUP COSTS</i>	77
TABLE 10 — <i>MONTHLY RECURRING COSTS</i>	77
TABLE 11 — <i>SAMPLE POLICY WITH SA FACTORS</i>	85
TABLE 12 — <i>SAMPLE SA POLICY WITH RA FACTORS</i>	86

1 Introduction

1.1 Document

1.1.1 Purpose

This Final Scientific and Technical Report summarizes the Concurrent Information Assurance Architecture (C-IAA) project and documents the accomplishments achieved under Air Force Research Laboratory (AFRL) contract F30602-01-C-0061 in accordance with its associated Statement of Work (SoW) and Contract Data Requirements List (CDRL) CLIN 0001, Item A005.

1.1.2 Overview

This document is organized into four primary sections: *Introduction* presents an overview of the C-IAA project; *Accomplishments* reviews tasks and products, findings, and recommendations; *Process* documents the conduct of the project; and *Research Products* presents the salient project results.

1.2 Project

1.2.1 Background

As organizations establish their presence on the Internet, and computer-based communication is increasingly used between members of an organization as well as with business partners, security becomes the forefront issue.

Currently, many tasks in the information security field are accomplished in a sequential manner, often after the fact, which limits the urgency of time response and usefulness of the tools and approaches currently available. The next step toward more secure networking is taking a Concurrent Information Assurance (C-IA) approach, which executes security-critical functionality concurrently on several different levels. The C-IA Architecture (C-IAA) postulates concurrent information assurance (IA) by providing configurable, coordinated, automated situation analysis, decision assistance, and response

Distributed, concurrent automated decision-making guided by configurable policy is a necessary ingredient for assuring IA in today's environment. The current state of the art allows for collection of data that may indicate attack. Additional research must be performed to determine the procedures for situation assessment, risk analysis, and response. Since the environment is highly dynamic as networks change topology and new exploits are discovered, the decision-making procedure must be configurable. The

speed at which input data is received and the complexity of that data necessitate automated assistance in decision-making. C-IAA would build on the existing technology and lead into the next paradigm level for enforcing network security.

1.2.2 Objectives

Architecture. The C-IAA's objective is to create the underpinnings for an architecture that executes security-critical functionality in an automated fashion, distributively, concurrently, and separately from other applications. C-IAA systems exploit the severability of concurrent processing into separate execution environments to achieve a high confidence and minimal impact on information, IA components, and the organizations dependant on that information.

Project. This project's purpose is to construct a hierarchy of intrusion detection (ID) and automated decision assistance modules that provide distributed and secure command and control of IA assets.

Scope. This effort's scope is to research, define, evaluate, and suggest the C-IAA's component design. Furthermore, this effort investigates and defines the attributes, requirements, costs, benefits, and practicality of the C-IAA and systems that conform to it.

2 Accomplishments

2.1 Tasks and Products

2.1.1 C-IA Architecture and Technologies

Research was conducted into technologies that exploit concurrent architectures for enforcing security, forensics evidence collection, ID and reaction, system survivability, and other IA research areas. Products of this basic research to establish the current state of the art for IA, which may be used to define C-IA as well as future OversightNets, OverNets, and OuterNets, can be found in this report as follows:

- C-IA Architecture, Section 3.4
- Secure Communications: State of the Art, Section 3.2
- Intrusion Detection, Analysis, and Response: State of the Art, Section 3.3
- DARPA Projects Relevant to C-IAA, Section 3.1

Automated security response was investigated at all levels from the processing node to the OuterNet, with results documented in Section 3.9 (Decision-making in C-IAA). Special consideration was given in that investigation to support security policies that include automated response as well human decision input.

Information assistance tools were researched to define an experimental toolkit for C-IA support. The collected toolkit was documented and delivered in accordance with CLIN 0001, CDRL A004 under separate cover.

2.1.2 OversightNet Secure Kernel

The secure isolation of C-IA critical functionality for the separation of application and security processing was investigated to determine the benefits of isolating each type of function within its own execution environment. Alternative approaches were considered to allow identification of an optimum C-IA OSK, with special consideration given to preserving the operability of currently used application software. The results of these studies are documented in Section 3.5 (C-IAA Processing Node).

2.1.3 OverNet

Requirements and design alternatives for the C-IA OverNet were researched and defined and documented in Section 3.7 (C-IAA OverNet). Investigations of interoperability issues, secure communications, and user authentication were performed and are addressed in Sections 3.2.4, 3.2, and 3.2.3, respectively.

2.1.4 OuterNet

Possible design approaches for the OuterNet were researched and defined, with special attention given to the issues associated with implementing the OuterNet as a physically separate network containing other networks. The results of these investigations are presented in Section 3.8 (C-IAA OuterNet).

2.2 Findings

The following findings are a distillation of the most salient outcomes of this project.

1. The postulated C-IAA was reviewed for overall usefulness and feasibility and was found to be an excellent architectural framework for the exploration of concurrent IA operations, as well as several other command and control technologies.
2. Much of the basic technology necessary to commence implementing a C-IAA compliant system is available today, although within a wide range of maturities. In many instances, initial prototype or experimental pilot implementations can be achieved using off-the-shelf products in conjunction with custom glue code.
3. Sufficient secure communications and authentication products exist to support the near- and medium-term needs of a C-IAA implementation.

4. The greatest obstacle to the envisioned implementation of
 - OversightNet is meaningful automation of IA analysis and administration functions.
 - OverNet is the non-technical barriers imposed by interorganizational relationships.
 - OuterNet is the non-technical political and legal constraints in the public environment.
5. A great deal of research into the broad area of knowledge engineering will be required before C-IAA can reach its full potential.
6. C-IAA, as for all research areas, must assess how functionality is achieved today. The next step desired must be compatible with existing practices to be used efficiently; if not compatible, it must offer a significant advantage over the existing tools and methods to justify the expense of new technology and training.

2.3 Recommendations

As the result of this research and consideration of C-IAA, WetStone makes the following recommendations:

1. Implementation of C-IAA components should be approached incrementally and experimentally, so that the effort can benefit from iteration, discovery, and the developments of compatible projects.
2. Of the recommended additional research efforts, first consideration should be given to those efforts that develop guidelines and overall system boundaries as C-IAA moves toward implementation.
3. Careful consideration should be given to the prioritization of C-IAA requirements. Although many functions are well within the state of the art, a few are considerably beyond, while others present a poor cost-benefit ratio on first analysis. The architecture should not be jeopardized by the premature inclusion of leading-edge capabilities.
4. Deliberate and steady growth of C-IAA capabilities from the processing node up toward the OuterNet is advised. Many of the proposed capabilities have not been previously integrated for these purposes. Therefore, unexpected advantages and obstacles during the course of development should be anticipated.

3 Research Products

3.1 DARPA Projects Relevant to C-IAA

The Defense Advanced Research Projects Agency (DARPA) has IA and Survivability programs that address solutions to national-level problems and high-risk, high-payoff technology development and exploitation such as sensor information technology.

There are several existing DARPA projects that contain concepts and/or implementations that can be used in C-IAA, with modifications or not. An examination of DARPA-sponsored projects was undertaken to identify projects that could benefit from, or be used in, the C-IAA. All currently funded DARPA projects were examined based on the project descriptions posted on the DARPA Web page [DARPA]. A more complete review would entail examining various projects' literature, briefings, and personal contact. This report can be used as the initial list of projects to be more closely examined.

Areas of research that can be tied into the C-IAA implementation include

- The control of agent-based systems, which allows rapidly assembling a set of disparate information systems into a coherently interoperating whole. This research can be tied into OversightNet, where sensors and networks rapidly change configuration, and into OverNet research, where OversightNets rapidly assemble and disassemble.
- The command post of the future, which allows the visualization and interaction environments needed for decision-making. Visualization and reaction components of decision engines should be a part of the command post of the future.
- Mobile networking technologies, which allow mobile nodes and sensors. This research can be tied into OversightNet, where sensors and networks rapidly change configuration, and into OverNet research, where OversightNets rapidly assemble and disassemble.
- Combined manned and unmanned operations, which allow for networks of intelligent unmanned warfare objects.
- Near-real-time planning and replanning, which allows better detection, correlation, and understanding of asymmetric threats (i.e. threats that arise from highly distributed, unconventional attackers).
- Advanced information technologies for logistics, which provide logistics command and control to the warfighter.

DARPA is divided into several offices, and each office sponsors one or more programs. The offices that have programs of interest to C-IA are

- DARPA Advanced Technology Office (ATO)
- DARPA Defense Sciences Office (DSO)

- DARPA Information Systems Office (ISO)
- DARPA Information Technology Office (ITO).

Table 1 lists DARPA projects that were found to be seeking functionality similar to functions required by C-IAA.

Table 1 — DARPA Projects Seeking Compatible Functionality

Program	Description	CIAA fit
ITO, Composable High-Assurance Trusted Systems	Secure operation of core network components	Secure oversight kernel research and implementation
ITO, Sensor Information Technology	Fuse sensor information	Situation and risk analysis and decision-making
ITO, Information Management	Rapidly acquire and manage massive amount of information	“
ITO, Translingual Information Detection, Extraction and Summarization	Find and interpret information from networked sources, regardless of language	“
ITO, Rapid Knowledge Formation	Enable knowledge experts to build knowledge databases without knowledge engineers as intermediaries	“
ISO, Evidence Extraction and Link Discovery	Knowledge discovery and data mining	“
ITO, Network Modeling and Simulation	Network modeling and simulation	Risk analysis
ATO, Command Post of the Future	Visualization support for decision-making	Visualization support for decision-making
ITO, Fault-Tolerant Networks	Automated response	Determining possible courses of action (CoAs)
ITO, Mixed Initiative Control of Autonomous Systems	Coordinate multilevel planning of distributed, semiautonomous systems	OverNet decision-making policy
ITO, The Joint Force Air Component Commander	Develop agile and stable air control in rapidly changing environment	OverNet and OversightNet decision making policy
ITO, Software Enabled Control	Control of autonomous objects	Decision-making policies
ITO, Active Networks	Network with “smart” packets capable of making decisions; studies in protocol interaction	OversightNet, OverNet interaction
ITO, Dynamic Coalitions	Dynamic collaboration of distributed agents	“
ITO, Autonomous Negotiating Teams	Decentralized, autonomous negotiation of roles and tasks, in real time.	“

Program	Description	CIAA fit
ITO, Dynamic Assembly for System Adaptability, Dependability, and Assurance	Software systems that can evolve and reconfigure dynamically	"

Additional detail concerning these projects is presented in the following sections. Relevant projects are identified within each DARPA office and programs supported within that office. Frequently, there are several relevant projects per program. Therefore, the overall program is described first, using the following format:

- DARPA office that sponsors the project; research area within the office; program name
- Description of the program with emphasis on how it relates to C-IAA
- Description of relevant projects

Each project is described using the following format:

- Project name
- Organization that performs the project
- Project description, with emphasis on how it relates to C-IAA
- Relationship to C-IAA, with suggestions on how the project can be used in C-IAA

3.1.1 Program: TIDES

ITO, Intelligent Software, Translingual Information Detection, Extraction, and Summarization (TIDES)

<http://www.darpa.mil/ito/research/tides/index.html>

Find and interpret information from networked sources, regardless of language.

Project: *Coalition TIDES: Machine Translation and Translingual Question Answering*

Organization: MIT Lincoln Laboratory, <http://www.ll.mit.edu/IST/>

Functionality: Information system that allows English-speaking users to access information in other languages, such as Korean news.

C-IAA Use: Use the accomplishments of this project to gain access to information sources in other languages that can provide information needed for situation assessment.

3.1.2 Program: RKF

ITO, Intelligent Software, Rapid Knowledge Formation (RKF)

<http://www.darpa.mil/ito/research/rkf/projlist.html>

Enables distributed teams of subject matter experts to build knowledge databases without knowledge engineers as intermediaries.

All projects from this program are relevant to C-IAA. A sample project is listed below.

Project: *KRAKEN: Knowledge Rich Acquisition of Knowledge from Experts Who Are Non-logicians*

Organization: CYCORP, Inc.

Functionality: Software to help subject matter experts create knowledge database without computer science expertise.

C-IAA Use: Use the accomplishments of this project to build depositories of knowledge, such as case studies and policy rules, to be used for situation assessment and risk analysis.

3.1.3 Program: Information Management

ITO, Intelligent Software, Information Management (IM)

<http://www.darpa.mil/ito/research/im/projlist.html>

Develop leading-edge technology to rapidly acquire, manage, exchange, and understand the massive amount of information relevant to the situation. Use these concepts to enhance C-IAA and allow retrieval of data necessary for automated decision-making, situation assessment, and risk analysis.

All projects are relevant to C-IAA. Sample projects are listed below.

Project: *Evolving Software Repositories*

Organization: National Institute of Standards, <http://netlib.org/utk/projects/esr>

Functionality: Access to public software repositories and mobile digital libraries.

C-IAA Use: Use the accomplishments of this project to build software repositories that could be useful for C-IAA. Applications are Domain Name Service (DNS) repositories or yellow pages.

Project: *Searching Unfamiliar Metadata*

Organization: University of California at Berkeley, <http://www.sims.berkeley.edu/research/metadata>

Functionality: Intelligent software that supports Entry Vocabulary for searching unknown data repositories.

C-IAA Use: Use these concepts to help search databases of known attacks and scenarios, case studies, risk analysis, etc.

Project: *Finding Information in Networked Environment*

Organization: Rutgers University, <http://aplab.rutgers.edu/ant>

Functionality: Intelligent software that supports finding information on WWW.

C-IAA Use: Use these concepts to help search databases of known attacks and scenarios, case studies, risk analysis, etc.

3.1.4 Program: Software Enabled Control

ITO, Embedded and Autonomous Systems, Software Enabled Control (SEC)

<http://www.darpa.mil/ito/research/sec/index.html>

This program deals mostly with controls of objects, such as autonomous vehicles. The issues of trusting the software pertain to C-IAA decision-making software.

Project: *Trustworthy Software: When Computers Serve as Proxies for Humans*

Organization: Oregon Graduate Institute,

<http://www.cse.ogi.edu/PacSoft/projects/TW/Default.htm>

Functionality: Methods and software to ensure trustworthiness of executable code.

C-IAA Use: Apply to C-IA software to ensure that automated decision-making and consequent actions preserve trustworthiness of C-IAA system.

Project: *Integrated Design and Analysis Tools for Software-based Control Systems*

Organization: University of California at Berkeley

Functionality: Software system based on discrete and continuous signals.

C-IAA Use: Exploit the experience in transitioning technology into production modality.

3.1.5 Program: MICA

ITO, Embedded and Autonomous Systems,

Mixed Initiative Control of Autonomous Systems (MICA)

Coordinate multilevel planning of distributed, semiautonomous forces with collective objectives.

No projects yet listed, just Broad Area Announcement (BAA).

3.1.6 Program: JFACC

ITO, Embedded and Autonomous Systems,

The Joint Force Air Component Commander (JFACC)

<http://www.darpa.mil/ito/research/jfacc/index.html>

Develop agile and stable air control in rapidly changing environment.

No projects yet listed.

3.1.7 Program: DASADA

ITO, Embedded and Autonomous Systems,
Dynamic Assembly for System Adaptability, Dependability, and Assurance (DASADA)
<http://www.darpa.mil/ito/research/dasada/index.html>

Software systems that can evolve and reconfigure dynamically.

Use these concepts to apply to C-IAA so that OversightNets and OverNets can have dynamic configuration.

No projects yet listed.

3.1.8 Program: ANTS

ITO, Embedded and Autonomous Systems,
Autonomous Negotiating Teams (ANTS)
<http://www.darpa.mil/ito/research/ants/index.html>

Highly decentralized and autonomous negotiation of tasks, roles, and allocations, in real time. Used mostly for operation of uninhabited vehicles.

Utilize these strategies for negotiation of policy between OverNet and OversightNet modules.

Project: *Planning Real-Time Negotiation for Mission-Critical Applications*

Organization: Honeywell Technology Center,
<http://www.htc.honeywell.com/projects/sa-circa/>

Functionality: Cooperation between uninhabited aerial vehicles.

C-IAA Use: Automatically create reactive intrusion response plans for autonomic computer security.

Project: *Autonomous Negotiating Teams and Model-Integrated Computing for Autonomic Logistics*

Organization: Vanderbilt University,
<http://www.isis.vanderbilt.edu/Projects/micants/micants.htm>

Functionality: Developed distributed problem-solving environments and negotiation algorithms for software components, with focus on maintenance operations.

C-IAA Use: Negotiation policies; technology transfer using Boeing experience.

3.1.9 Program: Active Networks

ITO, Networking and Distributed Systems, Active Networks
<http://www.darpa.mil/ito/research/anets/index.html>

“Smart” network packets, able to make autonomous decisions.

Project: *Virtual InterNetwork Testbed (VINT)*

Organization: Lawrence Berkeley National Laboratory, <http://netweb.usc.edu/vint/>

Functionality: Study in protocol scaling and interaction.

C-IAA Use: Application to OverNet-OversightNet communication protocols, especially with a hierarchy of OverNets.

3.1.10 Program: CHATS

ITO, Networking and Distributed Systems,
Composable High-Assurance Trusted Systems (CHATS)
<http://www.darpa.mil/ito/research/chats/projlist.html>

Provide mechanisms for secure operation of core network components.

Used for C-IAA OversightNet Secure Kernel.

Project: *Enhancing ReiserFS in Linux*

Organization: Namesys,
<http://www.namesys.com/v4.html>

Functionality: Enhance Linux security.

C-IAA Use: OSK.

Project: *Analyzing Security Policies for SE Linux*

Organization: Naval Research Laboratory, <http://chacs.nrl.navy.mil/projects/selinux/>

Functionality: Formal verification of existing and custom SE Linux security policies.

C-IAA Use: OSK.

Project: *Security Enhanced Bootstrap for Operating Systems (SEBOS)*

Organization: University of Maryland, College Park, <http://www.missl.cs.umd.edu/sebos/>

Functionality: Secure bootstrapping process.

C-IAA Use: OSK.

Project: *High-Assurance Open-Source Certificate Management System (CMS)*

Organization: BBN Technologies

Functionality: Open-source extension of BBN’s existing certificate management system (CMS) for X509 certificates.

C-IAA Use: Authentication between OverNets and OversightNets.

Project: *Secure File Sharing over the Internet Using SFS*

Organization: New York University, <http://www.scs.cs.nyu.edu/DARPA/sfs>

Functionality: Enable secure networking using open-source operating systems.

C-IAA Use: OSK.

Project: *Secure Auditing for the Linux Kernel*

Organization: SPAWAR Systems Center, San Diego, <http://secureaudit.sourceforge.net>

Functionality: Software that audits Red Hat Linux OS.

C-IAA Use: OSK, forensic tools.

Project: *Portable Open-Source Security Elements (POSSE)*

Organization: University of Pennsylvania, <http://www.cis.upenn.edu/~posse>

Functionality: Enhance security of BSD-licensed open-source operating systems, OpenBSD.

C-IAA Use: OSK.

3.1.11 Program: Dynamic Coalitions

ITO, Networking and Distributed Systems,
Dynamic Coalitions

<http://www.darpa.mil/ito/research/dc/index.html>

Technologies for dynamic collaboration of distributed agents.

Used for C-IAA collaboration between OversightNets and OverNets, as well as OverNet hierarchy.

Project: *MSME: Multidimensional Security Management and Enforcement*

Organization: BBN Technologies

Functionality: Framework and tools for defining and negotiating security policy between coalition members.

C-IAA Use: Collaboration between OversightNets and OverNets, and OverNet hierarchy.

Project: *Methodologies for Reliable Certificate Revocation*

Organization: Drexel University

Functionality: Detecting coalition members that turn hostile.

C-IAA Use: Collaboration between OversightNets and OverNets, and OverNet hierarchy.

Project: *High-Performance, Robust, and Secure Group Communication for DC*

Organization: Johns Hopkins University, http://www.cnds.jhu.edu/funding/dynamic_coalitions/

Functionality: Scalable secure protocols based on SPREAD group communication system (<http://www.spread.org>) and the CLIQUES key agreement protocol suite (<http://www.isi.edu/~gts/CLIQUES/>).

C-IAA Use: Collaboration between OversightNets and OverNets, and OverNet hierarchy.

Project: *Specifying and Enforcing Security Policies in Multiparty Communication Systems (Antigone 2.0)*

Organization: University of Michigan, <http://antigone.citi.umich.edu>

Functionality: Software tool that lets users specify security policy and automatically configures protocols that implement the policy. Policy is based on several dimensions and variable security needs.

C-IAA Use: Collaboration between OversightNets and OverNets, and OverNet hierarchy.

Project: *Flexible Coalition Policies for Secure Information Sharing*

Organization: Veridian-PSR

Functionality: Risk assessment of sharing information between coalition partners.

C-IAA Use: Risk assessment.

3.1.12 Program: Fault-Tolerant Networks

ITO, Networking and Distributed Systems,
Fault-Tolerant Networks (FTN)
<http://www.darpa.mil/ito/research/ftn/index.html>

Project: *Active Network Intrusion Detection Response*

Organization: NAI Labs

Functionality: Framework and tools that allow network users to reprogram and customize routers, firewalls, switches, and other components to provide new network services on the fly. Based on the Intruder Detection and Isolation Protocol. IDIP provides cooperation among ID systems, firewalls, routers, network management components, and hosts so that intrusions that cross multiple network boundaries can be automatically traced and blocked as close to their sources as possible.

C-IAA Use: Situation assessment, automated response.

3.1.13 Program: NMS

ITO, Networking and Distributed Systems,
Network Modeling and Simulation (NMS)
<http://www.darpa.mil/ito/research/ftn/index.html>

Technologies for network modeling and simulation.

Used for C-IAA situation assessment and risk analysis.

Project: *Maya: Next-Generation Performance Prediction Tool for Global Networks*

Organization: University of California, Los Angeles (UCLA)

Functionality: Framework for evaluation of network performance, including fault localization and recovery.

C-IAA Use: Situation assessment and risk analysis.

3.1.14 Program: Sensor Information Technology

ITO, Networking and Distributed Systems,
Sensor Information Technology

<http://www.darpa.mil/ito/research/sensit/index.html>

Networks of sensors of various types: micro, wireless, etc. Fusion of information.

Output to be used in OversightNet decision-making.

Project: *Reactive Sensor Networks*

Organization: The Pennsylvania State University, <http://strange.arl.psu.edu/RSN/>

Functionality: Collaborative signal processing software has been implemented for Windows and Linux.

C-IAA Use: OversightNet information collection.

3.1.15 Program: EELD

ISO, Evidence Extraction and Link Discovery (EELD)

<http://dtsn.darpa.mil/iso/index2.asp?mode=9>

Used for C-IAA situation assessment and risk analysis, and as a part of forensic utilities.

No projects yet listed, just the BAA.

3.1.16 Program: Command Post of the Future

ATO, Command Post of the Future

<http://www.darpa.mil/ato/programs/cpof.htm>

Provides visualization support for situation and risk analysis.

Use in C-IAA for situation and risk analysis.

3.1.17 Program: Cyber Panel

ATO, Cyber panel

<http://www.darpa.mil/ato/programs/cyberpanel.htm>

No information provided.

3.1.18 Program: Dynamic Coalitions

ATO, Dynamic Coalitions

No information provided.

3.1.19 Program: Fault-Tolerant Networks

ATO, Fault-Tolerant Networks

No information provided.

3.2 Secure Communications: State of the Art

The C-IAA consists of a collection of hierarchical, cooperating processes, which may be physically distributed across multiple physical platforms and in turn geographically distributed. When remote C-IAA processes exchange IA data or command and control (C²) information, such communications must be immune to subversion, or C-IAA itself would be vulnerable to the very attacks it exists to prevent and control. In this regard, C-IAA has the same need for secure communications as any other C² architecture and shares the same basic requirements for

- privacy and integrity of messages in transit, and
- authentication of correspondents.

Contemporary demands for ease of integration, mobile correspondents, and rapid deployment further complicate the security environment.

3.2.1 Virtual Private Networks

A VPN allows distributed private networks to communicate securely with each other over untrusted public networks. According to Infonetics, in 1999, corporations bought US\$281 million of VPN equipment [Youn00]. Multinational corporations and complex industries are the prime candidates for VPN use. One company executive claims that “the total cost of ownership of a VPN internet connection is so much less than one using dedicated lines – on the order of 30% to 70% - that we contend no company can ignore the technology” [SecuVPN].

The typical Internet Protocol (IP)-based VPN transfers packets that have been encrypted and authenticated between two VPN nodes by encapsulating the packets underneath a special VPN protocol header, thereby creating a so-called secure *tunnel* between the VPN nodes. This operation is transparent to users of the VPN, thus correspondents perceive a private network between their remote subnetworks. When more than two subnetworks participate in a VPN, a virtual backbone is constructed of secure tunnels between the VPN nodes. Transport for the protected, encapsulated packets may be any of the variety of methods that support IP on or off the Internet.

In general, the process of sending a message using a VPN is as follows:

- A host that participates in a VPN sends clear traffic to a VPN edge device located at the point of connection to the public network.
- The edge device examines the data according to rules specified by the network manager, securing the information or allowing it to pass unaffected.
- When data protection is required, the edge device encrypts and authenticates the packet.

- The edge device then prepends a new VPN header to the resulting protected data and transmits the new packet as a standard IP packet.

Upon receipt, the corresponding edge device “unwraps” the encapsulated packet and sends the resulting original packet to the destination host.

VPN technology is mature and has been reduced to practice by a variety of commercial, off-the-shelf (COTS) products and systems. Incorporation of a VPN into a larger network design requires trade-off and selection among the leading implementation approaches:

- Router versus firewall extensions¹, and
- Software versus hardware cryptography to support encryption and authentication.

Furthermore, trade-off and selection of the following implementation specifics must be made:

- *Encryption algorithm* – such as DES/3XDES, RSA, RC4, RC5, or IDEA.
- *Key Exchange Protocol* – such as IKE or SKIP.
- *Tunneling Protocol* – such as Layer 2 Tunneling Protocol (L2TP), Internet Protocol Security (IPSec), PPTP², CIPE, or PPP with SSH.
- *Certification process* – such as dedicated Certification Authorities³ (CAs), third-party CAs, or alternatives to certificates such as S-expressions.
- *Authentication process* – using passwords, soft tokens, or hard tokens.
- *Usage policy* – permitting data transport through the VPN.

3.2.2 Secure Transport Protocols

3.2.2.1 IPSec and IPv6 Protocols

IPSec is a network-layer protocol, designed to ensure security in IP-based communication systems. IPSec establishes host-to-host security by permitting two systems to establish an encrypted TCP session. IP payloads are encrypted and encapsulated in an IP header

¹ Most commercial firewalls, such as Secure Computing's *Sidewinder* [SecuVPN], optionally support VPN functions.

² PPTP is a Microsoft protocol for VPN [RFC1171] that is an extension of PPP and supported under Linux, but is known to have serious security issues. CIPE is a kernel-level network encryption system that is considered to be best suited to enterprise setups. The default protocol and key exchange algorithm in commercial VPN modules today are IPSec and IKE. The default protocol for ISP use is L2TP. All major firewalls support IPSec.

³ Many vendors that offer commercial VPNs act as their own top-level CA and issue certificates for their customers, such as Kolumbus/SSH [Kolu01].

for secure transfer over the Internet. IPSec is implemented as a software library compatible with IPv4. IPSec is a part of the IPv6 standard.

IPSec supports the key management requirements of network layer security by using the Internet Key Management Protocol. IKMP is an application-layer protocol that is independent of the lower-layer security protocol. IKMP is based on the ISAKMP/Oakley. Internet Key Exchange (IKE) protocol [RFC2409] is used to exchange the keys. Current IPSec protocols and algorithms [RFCs 2401-2412, 2085, 2104, and 2451] can exchange keying material using IKE [RFC2409] and protect dataflows using the IP Authentication Header (AH) [RFC2402] protocol and/or IP Encapsulating Security Payload (ESP) protocol [RFC2406].

The protocol formats for AH and ESP are independent of the cryptographic algorithm. IPSec flexibly supports combinations of authentication, integrity, access control, and confidentiality. IPSec can be configured to use DES, 3DES, RC5, Cast, Idea, or Blowfish encryption algorithms. Because single DES has been broken in less than 23 hours, IPSec establishes a new encryption key every X hours or Y bytes, where X and Y are user-configurable.

ISAKMP/Oakley is a heavyweight protocol that requires hosts to exchange several packets to set up a “security association.” Therefore, the initial set up cost is high, and this protocol is most suitable for reliable connections between few correspondents.

The other key distribution protocol that can be used for IPv6 is SKIP. SKIP is a lightweight protocol, where encryption keys are securely hidden within a packet. If secure communication fails, there is no fallback. Therefore, SKIP is most suitable for numerous, short-term connections.

IPSec flexibility creates complex management tasks that become especially difficult as networks scale up and require different security policies, controlled by different entities, for different kinds of traffic in different parts of the network. A basic feature of IPSec is that two hosts can establish a security association even though they might not share a common security policy or trust one another at all.

The IP Security Policy (IPSP) Working Group created the IPSec Policy Protocol, configuration policy model, and management information base to provide a scalable, decentralized framework for managing, discovering, and negotiating the host and network IPSec policies that govern access, authorization, cryptographic mechanisms, confidentiality, data integrity, and other IPSec properties.

For example, the security policy could be rendered as a(n):

- Lightweight Directory Access Protocol (LDAP) [LDAP] schema in a directory

- On-the-wire representation over a transport protocol like the Common Object Policy Service (COPS) [COPS, COPSPPR]
- Text-based policy specification language suitable for editing by an administrator
- Extensible Markup Language (XML) document

IPSec is the major VPN protocol used for business applications. The Automotive Network Exchange (ANX) VPN is one of the largest operational VPNs, connecting automotive manufacturers and their suppliers.

3.2.2.2 SSL, TLS, and S-HTTP Protocols

Secure Sockets Layer v2 and v3 (SSL) and Transport Layer Security v1.0 (TLS) protocols are client-server-oriented protocols designed to provide privacy and authentication between two applications communicating over TCP. SSL and TLS encrypt traffic between two specific ports over TCP/IP.

SSL was developed by Netscape [SSL96]. TLS is the successor protocol to SSL [TLS]. TLS and SSL are widely used with the HTTP protocol, for example, when supplying payment information securely in e-commerce applications. Many vendors sell commercial versions of SSL and TLS, such as SpyruS [Ther01]. TLS is also being used for adding security to many other common protocols that run over TCP, such as SMTP [RFC2487].

SSL is application-protocol-independent and is composed of two sublayers. The top sublayer, SSL Handshake Protocol, allows the would-be correspondents (e.g., client and server) to negotiate an encryption algorithm and cryptographic keys before the application transmits or receives its first byte of data. The lower sublayer, SSL Record Protocol, is used for encapsulating higher-level protocols for transmission over reliable channels such as TCP/IP.

In SSL, encryption is used after an initial handshake to define a secret key. Symmetric cryptography is used for data encryption (e.g., DES [DES], RC4 [RC4], etc.) The peer's identity can be authenticated using asymmetric, or public key, cryptography (e.g., RSA [RSA], DSS [DSS], etc.). Message transport includes a message integrity check using a keyed MAC. Secure hash functions (e.g., SHA, MD5, etc.) are used for MAC computations.

SSL is designed to ensure interoperability, so that independent applications conforming to SSL 3.0 can successfully exchange cryptographic parameters without knowledge of one another's code. Since the cryptographic and other algorithms are negotiated, new public key and bulk encryption methods can be incorporated as necessary.

Secure Hyper Text Transfer (S-HTTP) protocol is an alternative to SSL and TLS, specifically for use on WWW, as an enhancement to HTTP traffic. It uses RSA public-key encryption, and is supported by America Online, CompuServe, IBM, Netscape, Prodigy, SPRY, and Spyglass.

3.2.2.3 Secure Shell Protocol

Secure Shell (SSH) provides support for secure remote login, secure file transfer, and secure TCP/IP and X11 forwarding on Unix systems. SSH is implemented at the application level and operates over TCP/IP or other reliable but insecure transports. An IETF Working Group is attempting to enhance SSH so that it provides strong security and works “reasonably well” without a global certification infrastructure [Secsh01]. SSH is the only protocol for secure remote X-windows connections.

3.2.2.4 Socket-S Protocol

Socket-S (SOCKS) is networking proxy protocol that enables hosts on one side of a SOCKS server to gain full access to hosts on the other side of the server without requiring direct IP reachability. SOCKS redirects connection requests from hosts on opposite sides of a SOCKS server. The SOCKS server authenticates and authorizes the requests, establishes a proxy connection, and relays data. SOCKS is commonly used as a network firewall that enables hosts behind a SOCKS server to gain full access to the Internet, while preventing unauthorized access from the Internet to the internal hosts.

There are two major versions of SOCKS: SOCKS v4 and SOCKS v5, both publicly available. Both packages include clients for telnet, FTP, finger, and whois. The SOCKS v5 Reference Implementation also includes archie, ping, and traceroute. Many commercial products include built-in SOCKS protocol support, such as Permeo’s e-Border product family [SOCKSF].

3.2.2.5 IDXP and IDMEF Protocols

The ID Exchange Protocol (IDXP) provides for the exchange of ID Message Exchange Format (IDMEF) messages, unstructured text, and binary data between ID entities. IDXP uses the BEEP framework, described in Section 3.2.3.3.3.

IDMEF [IDMEF] defines data formats and exchange procedures for sharing information of interest to ID and response systems and to the management systems that may need to interact with them. IDMEF is written in XML and is oriented toward network ID.

3.2.2.6 Spread

The Spread Wide Area Group Communication System [SPREAD] was developed under DARPA funding to provide wide-area secure communications. Spread integrates two

low-level protocols: one for local area networks called Ring, and one for the wide area network connecting them, called Hop. A Spread implementation is publicly available on the Web.

3.2.2.7 L2TP Protocol

Currently, dial-in users use Point-to-Point Protocol (PPP) over Serial Line IP (SLIP). L2TP protocol was proposed by the IETF to create and maintain VPNs over public TCP/IP connections using PPP. L2TP encapsulates PPP frames to be sent over IP, X.25, frame relay, or ATM. L2TP is a network-layer protocol and is used by Internet Service Providers (ISPs) to provide remote dial-up VPN access.

Current implementations of L2TP include Microsoft's MS-PPTP and Microsoft's and Cisco's joint L2TP. Both have been submitted for standardization and are presently in many commercial VPN products such as Aventail and Freegate, for Linux and Windows NT distributions. Comprehensive security analysis of MS-PPTP showed vulnerability to attacks, such as offline password guessing with tools such as L0phtcrack [Opp100].

Since L2TP uses PPP for encapsulation, it does not require installation of an extra package on the remote client.

3.2.3 Authentication Technologies

Authentication is a necessary component for secure communications. The strongest forms of authentication involve the combination of more than one authentication factor:

- *What you know* – Passwords, passphrases.
- *What you possess* – Physical tokens, private keys.
- *What you are* – Biometrics.

3.2.3.1 Password-based Authentication

Password-based authentication is the simplest and the least secure. There are several alternative ways to use passwords securely:

- Weak password authentication uses plaintext passwords.
- Medium-strong password authentication methods include SSH password authentication and key exchange, password-protected TSL/SSL certificates, and default password pre-authentication in Kerberos V.
- Strong password authentication methods include SRP, EKE, SPEKE, OKE, AMP, and numerous other protocols used for e-commerce.

3.2.3.1.1 Secure Remote Password

Secure Remote Password (SRP) protocol [RFC2945] uses hashing for negotiating secure connections via a user-supplied password. This system also performs a secure key exchange in the process of authentication, allowing security layers (privacy and/or integrity protection) to be enabled during the session. Trusted key servers and certificate infrastructures are not required, and clients are not required to store or manage any long-term keys. An SRP implementation is publicly distributed. Some IPSec protocols use SRP.

3.2.3.1.2 Kerberos Authentication System

The Kerberos protocol is used as an alternative to one-time passwords. It is used to prove to a server that a client is running on behalf of a particular user. The client must first contact a separate, authentication server and initiate session key generation. Kerberos cannot handle digital signatures and many other useful features. Kerberos exchanges DES-encrypted messages among participants.

Kerberos installation requires modifications in client and server software but is bundled with many software packages. For example, Windows 2000 adopted Kerberos 5 with extensions for public key authentication as the default protocol for network authentication [Mic99].

3.2.3.2 Tokens, Keys, Certificates, Signatures

The following sections present various authentication schemes based on “what you have.” Tokens and keys are the most basic components, out of which signatures and certificates can be built. Various hardware and software devices and protocols are used for storing and authentication. For example, hardware tokens such as smart cards are usually relatively expensive (at least \$10 per user) and require hardware equipment and maintenance.

Public key cryptography is based on the mathematics of public and private keys. Public Key Infrastructures (PKIs) provide the means to authenticate users based on public key cryptography. The main problems in use are the secure storage and portability of private keys and the wide distribution of public keys in certificates by trusted authorities.

3.2.3.2.1 PKCS #15

PKCS #15 is a Cryptographic Token Information Format Standard by RSA Laboratories. It specifies the syntax for storing digital credentials (keys, certificates, etc.) on cryptographic tokens and how this information is to be accessed.

PKCS#15 is designed to ensure interoperability, so that users can employ cryptographic tokens to identify themselves to multiple, standards-aware applications, regardless of the application's crypto-key (or other token interface) provider.

3.2.3.2.2 Digital Signatures

Digital signatures enable authentication, accountability, and data integrity for electronic transactions and are currently used in a wide range of online business transactions. Using a PKI, digital signatures can, uniquely and unequivocally, identify an entity or individual [Mosok]. Typically, 128-bit hashes are used, such as SHA or MD5.

Digital signatures are produced in a two-step process. First, a mathematical function is performed on a message producing a unique hash code or digest. Then, the user encrypts the hash with his or her private key, producing a digital signature. This action cryptographically binds the signer to the message and can be used to prove that the signer was in possession of the message in the exact form that produced the hash. The original message and the digital signature can then be transmitted.

Once received, the recipient ensures authentication and verifies data integrity in a two-step process as well. First, the digital signature and hash function is decrypted using the sender's public key. The sender's public key will be the only key able to decrypt the encrypted hash. Second, the message is passed through the same hash function to verify data integrity. If the data has been modified in any way, the hash function will not produce the same result. With secure storage of the private key, the recipient can be sure that the message was indeed possessed by the sender.

Although digital signatures do not address data privacy, they have wide-scale acceptance, and most modern economies have existing or pending legislation giving digital signatures legal recognition [Entrust]. In the United States, the *Electronic Signatures in Global and National Commerce Act (E-Sign)* was passed in November 1999, making digital signatures legally binding. Entrust Technologies, VeriSign, and Arcot Systems are just a few of the many organizations with digital signature services or software.

3.2.3.2.3 Certificates

Certificates form a cornerstone of most PKI implementations, whether public or private, open or closed. The purpose of certificates is to allow two arbitrary parties to interact securely using public key cryptography without having a prior special arrangement for the exchange of their public keys. Instead, a Trusted Third Party (TTP), accepted by both corresponding parties, creates certificates that contain each party's name and public key. Because of special safeguards maintained by the TTP, which is also known as a CA, the certificates of the corresponding parties may be openly published and used as necessary. Relying parties accept the binding of the name to the public key within the certificate.

Then public key cryptography methods can be used for authentication and encryption. However, currently, certificates are most often used for authentication only.

The leading certificate format currently is X.509, version 3, which can convey a variety of information useful to the associated cryptographic processes, all protected and assured by the CA, including security policy, certificate use restrictions, expiration date, transaction value limits, subject's contact information, valid cryptographic algorithms, etc.

The IETF Public Key Infrastructure (PKIX) Working Group, W.509, developed a PKI for the Internet based on ITU-T recommendation X.509 [PKIX]. A family of RFC documents describes management of public key certificates and revocation lists. Commercial products are often PKIX compliant. WWW browsers support certificate hierarchies, validation, and storing of certificates locally. Sites that use PKI-compliant products must purchase server certificates from third parties like VeriSign or Thawte, or they must install and administer their own certification authorities.

There are many ways to use X.509 certificates. For example, TTPs could use globally or locally accepted certificates. There could be several certification hierarchy roots, and policies for their interaction and their spawning of lower-level certification authorities must be specified. Authentication domains can be treated as composable objects [HOSANA].

X.509 certificates do not address the global namespace problem, which leads to authorization problems, so the Simple Distributed Security Infrastructure (SDSI) was developed as an alternative to PKIX. SDSI uses S-expressions instead of X.509 certificates [SDSI].

ANSI developed attribute certificates to augment X.509 certificates, and these certificates may be used in TSL [Oppl00].

3.2.3.2.4 Smart Cards

Cryptographic smart cards are physical tokens that contain a CPU and enough memory to store private keys and perform cryptographic operations such as digital signatures. They are usually PIN-protected and offer some form of hardware-based tamper-resistance. Smart cards offer a wide range of security assurance levels, with higher assurance directly related to higher cost. Their adoption in the US has been extremely slow due to the requirement of deploying smart card readers within any using infrastructure and the expense of that deployment compared to competing methods available in the US. Ironically, smart cards have seen widespread adoption in Europe and Asia due to the lack of those competing methods.

3.2.3.2.5 Software Tokens: ArcotID

Arcot Systems [ArcotID] offers ArcotID, a software-based alternative to hardware authentication tokens. ArcotID uses a technique known as *cryptographic camouflage*. Arcot stores the private key, using the encryption key as the PIN. If an incorrect PIN is supplied, no key can be retrieved or several seemingly correct keys are retrieved. Arcot claims that the only way to find the correct key is to try out all retrieved keys, which would expose the attacker. An Arcot system is used to authenticate all electronic Visa purchases by requiring the user to supply authentication to the bank that issued the card.

3.2.3.2.6 PDM Protocol

The PDM protocol [PDM] enables a user to acquire cryptographic credentials, such as private keys and PKCS#15 structures, from a workstation with locally trusted software but with no user-specific configuration. Using PDM protocol is less secure than using a token, but may be useful until hardware tokens become ubiquitous or as a backup strategy when a user's token is lost or malfunctioning.

3.2.3.2.7 Privacy Enhanced Mail

Privacy Enhanced Mail (PEM) is an application-layer protocol used to enhance exchange of documents over e-mail. It applies encryption, source authentication, and integrity protection to e-mail messages, assuming a rigid CA hierarchy [RFC1421-1423].

3.2.3.2.8 Pretty Good Privacy (PGP)

PGP is an authentication mechanism, used for protecting e-mail and file storage by digitally signing and encrypting information "objects." It is well suited for any store and forward application. PGP assumes that each user, independently and at his own risk, decides which certificates to trust.

3.2.3.3 Authentication Mechanisms

There are several standardized Application Program Interfaces (APIs) that can be used to bring together the various authentication systems. An API supports a range of underlying mechanisms and technologies and hence allows source-level portability of applications to different environments. Possible authentication APIs include

- Generic Security Service Application Program Interface (GSS-API)
- Federated Naming Specification (XFN)
- BEEP
- RADIUS

3.2.3.3.1 GSS-API

GSS-API has been standardized as [RFC2078]. A typical GSS-API caller is itself a communications protocol, invoking GSS-API to protect its communications with authentication, integrity, and/or confidentiality security services. A GSS-API caller accepts tokens provided to it by its local GSS-API implementation and transfers the tokens to a peer on a remote system; that peer passes the received tokens to its local GSS-API implementation for processing. The security services available through GSS-API have been implemented over a range of underlying mechanisms based on secret key and public key cryptographic technologies.

3.2.3.3.2 XFN

The Open Group (X/Open) proposed a technical standard called Federated Naming Specification (XFN), which uses a federated naming service together with an API and specifies the naming policies to be used in conjunction with this service [XFN]. XFN provides a method for federating multiple naming services under a single, uniform interface for the basic naming operations. Applications use API, and the service decides what authentication system to invoke and how.

3.2.3.3.3 BEEP

Blocks Extensible Exchange Protocol Framework (BEEP) [BEEP, BEEP01] is an application protocol framework for connection-oriented, asynchronous request/response interactions. BEEP specifies initiating connections, framing, managing security, and multiplexing multiple channels in a single authenticated connection.

BEEP can be used to construct a VPN by creating an application-layer tunnel that transparently forwards data via a chain of proxies.

3.2.3.3.4 RADIUS

RADIUS is a vendor-independent protocol that allows multiple dial-in access points (through serial lines and modems) so that users can access a centralized user database. RADIUS works on a client-server model. Clients pass user connection requests to the server. The RADIUS server keeps authentication, authorization, and type-of-service configuration on each user (for example, should the user use SSL, SSH, PPP, telnet, or login). The server authenticates the user and returns all configuration information necessary for the client to deliver service to the user.

3.2.4 Interoperability through Standardization

Some issues in protocol development are standardization (e.g., is OAKLEY interpretation of IKE was standardized, SKIP interpretation is optional); proprietary versions (e.g.,

Microsoft includes a proprietary version of Kerberos in Windows 2000); global use (e.g., cryptographic export laws make some security software available only in the US); and interoperability.

Interoperability can be assessed on many levels, including between APIs, protocol layers, applications, gateways, networks, security levels, and many others. Individual technologies must be evaluated, as well as the various implementations offered by different vendors, both separately as well as working with other technologies; for example, are there any security holes when executing Kerberos over IPSec?

Evaluating protocols and their implementation is a mushrooming research area, and it usually falls into the category of (informal) protocol analysis and software testing and/or formal methods. For example, SSL and its derivative TLS protocols are widely used in e-commerce today for HTTP traffic and are distributed with Netscape and Internet Explorer. Some known vulnerabilities are that SSL works on TCP and not on UDP traffic and that it interacts poorly with proxy servers [Oppl00]. The Bleichenbacher attack [Blei98] is based on mathematical manipulation of cryptographic functions. [AlFo98] describes multiprotocol attacks that can be used to break otherwise secure public key-based authentication protocols. These attacks are possible when the public key infrastructure permits the use of a user's public key in multiple protocols.

Evaluating individual technologies is necessary to ensure interoperability, which is a key factor in ensuring widespread use of technology. Various standardization efforts are under way to ensure seamless deployment, such as the IETF efforts outlined in this section. Interoperability research is well on its way in the e-commerce world. The WWW Consortium and Commerce Net Consortium formed a Joint Electronic Payment Initiative (JEPI) to ensure that multiple paying schemes, protocols, and transport mechanisms interoperate. IBM performed research to unify different payment mechanisms in a common framework with corresponding APIs [PeAsStWa98].

In the current market with many vendors and many COTS and GOTS products, users do not have reliable means of ensuring product reliability for secure networked and distributed systems. A "Consumer Guide to Networked and Distributed Security Products" would be the first step, and a more rigorous evaluation, such as Trusted Computer Security Evaluation Criteria (TSEC), leading to certification would be ideal. So far, the only work done in this area has been performed by Trusted Network Interpretation of TSEC [TSEC]. A TSEC C2 certificate was granted to Windows NT [Oppl00]. In Europe, TSEC was followed by information technology evaluation criteria (ITSEC). The US, Canada, and Europe proposed a common criteria (CC) to the International Standards Organization (ISO) in December 1997 and have been evolving it since then [CC]. This work could possibly be extended toward networked and distributed systems. According to [Neum00], CC does not pose excessive system requirements such as TCSEC; rather, it serves as a framework for evaluation and distinguishes between

functional and assurance requirements. [Neum00] proposes a set of recommendations for evaluation criteria for survivable systems and networks.

Several organizations provide evaluation and certification services: namely, ICSA, NSA, and NAIP. ICSA labs provide test configurations to verify interoperability of IPSec vendor equipment [ICSA]. NSA provides security certification to commercial products and vendors, such as to Motorola's AIM chip [Signal98]. The National IA Partnership (NIAP) is a US Government initiative started in 1997 to meet the security testing, evaluation, and assessment needs of both information technology (IT) producers and consumers. NIAP is a collaboration between the National Institute of Standards and Technology (NIST) and the National Security Agency (NSA) in fulfilling their respective responsibilities under the *Computer Security Act of 1987*. The goal of the partnership is to "promote the development of technically sound security requirements for IT products and systems and appropriate metrics for evaluating those products and systems. The long-term goal of NIAP is to help increase the level of trust consumers have in their information systems and networks through the use of cost-effective security testing, evaluation, and assessment programs" [NIAP].

The first annual Federal IA Conference (FIAC), sponsored by the Federal Business Council, was held in October 2001. The inaugural conference, entitled "An Alliance for a More Secure Nation," is being designed specifically to meet the IA needs of the federal government and its workforce.

The IETF governs the adoption of Internet standards through a process that includes Requests for Comments (RFCs) and a wide variety of Working Groups that consider new standards [IETFWG]. The existence of a particular Working Group indicates wide interest in the topic for which it is responsible. Current security-related Working Groups are identified in Table 2.

Table 2 — Security-related IETF Working Groups

IETF Working Group	Goal	Output
Authenticated Firewall Traversal	Specify a traversal protocol supporting both TCP and UDP applications with a general framework for authentication of firewall traversal. To promote interoperability, the group proposed a base authentication technique for use within the general authentication framework.	SOCKS protocol and GSS-API for SOCKS

IETF Working Group	Goal	Output
Common Authentication Technology	Provide distributed security services (authentication, integrity and confidentiality, possibly authorization) to a variety of protocol callers in a manner that insulates those callers from the specifics of underlying security mechanisms.	GSS-API for C and Java
Intrusion Detection Exchange Format	Define data formats and exchange procedures for sharing information of interest to ID and response systems and to management systems that may need to interact with them.	IDEMF XML format, IDXP protocol
IP Security Protocol	Develop mechanisms to protect client protocols of IP, such as IP security architecture, secure transmission, key management protocol, IP security protocol.	IKMP, IKE, ISAKMP/Oakley, the IP AH, and IP ESP
IP Security Policy	Negotiation, exchange, storage, and specification language for IPSec policies.	IPSec configuration policy model, IPSP requirements, IPSec policy information base, IPSec policy configuration MIB
IP Security Remote Access	IPSec for remote access such as dial-in.	Dynamic Host Configuration Protocol (DHCP), PIC protocol, requirements for IPSec remote access scenarios
Kerberized Internet Negotiation of Keys	Define centralized key management as an alternative to IKE, using the Kerberos architecture for key management.	Kerberized Internet Negotiation of Keys (KINK) and requirements
Kerberos WG	Kerberos specification.	Kerberos V
Multicast Security	Specify secure group communication over Internet. Provide scalable solutions for groups with a single source and a very large number of recipients, where the data is transmitted via IP-layer multicast routing protocols (with or without guaranteed reliability). Each group has a single trusted entity (the Group Controller) that sets the security policy and controls group membership.	The group domain of interpretation, Group Secure Association Key Management Protocol (GSAKMP), Group Key Management Architecture (GKMA), GSAKMP light, group security policy token
An Open Specification for Pretty Good Privacy	Standardize PGP.	OpenPGP Message format, MIME security with OpenPGP

IETF Working Group	Goal	Output
One Time Password Authentication	Standardize one-time password technology, using the technology in the Bellcore S/KEY system and related interoperable packages (e.g., logdaemon, NRL OPIE).	OTP extended responses, a one-time password system, the one-time-password SASL mechanism
Public Key Infrastructure (X.509)	Develop Internet standards needed to support an X.509-based PKI.	X.509 version 3 certificates, CRLs version 2, the Certificate Management Protocol (CMP), the Online Certificate Status Protocol (OCSP), the Certificate Management Request Format (CRMF), second certificate management protocol (CMC), and many others
Securely Available Credentials	Provide credential portability.	PDM protocol, securely available credentials framework and protocol
Secure Shell	Update and standardize SSH.	SSH
S/MIME Mail Security	Specify security for S/MIME mail.	S/MIME certificates, cryptography, etc.
Secure Network Time Protocol	Define the message formats and protocols - specifically, modifications to the existing Network Time Protocol (NTP) - that are necessary to support the authenticated distribution of time for the Internet.	Public key cryptography for NTP version 2
Security Issues in Network Event Logging	Make syslog more secure. Syslog is a de facto standard for logging system events.	Syslog documentation, secure Syslog
Transport Layer Security	Transport layer security.	TLS protocol, HTTP over TLS, Kerberos in TLS
XML Digital Signatures.	Develop a simple, extensible XML digital signature syntax, i.e. XML-compliant syntax used for representing the signature of Web resources and portions of protocol messages (anything referred by a URL) and procedures for computing and verifying such signatures. Such signatures may be able to provide data integrity, authentication, and/or non-repudiability.	XML signature requirements, syntax, and processing; Canonical XML vs. 1.0

3.3 Intrusion Detection, Analysis, and Response: State of the Art

A brief overview of approaches currently used to detect attacks, assess current situations, and analyze potential risks is presented in the following subsections.

The attack detection approaches can be roughly divided into those that recognize known attack signatures and those that recognize behavior deviant from the “normal” behavior. In this report, we will focus on the approaches relevant to C-IAA, i.e. those approaches that use data obtained from ID tools to infer attack presence. Knowledge engineering techniques are used to “fuse” the data and assess possible attack presence.

3.3.1 Intrusion Detection

ID is one aspect of IA that deals with detecting inappropriate, incorrect, or anomalous activity. ID tools operate in two realms: on a host (host-based) and at the network level (network-based). In both approaches, the most common detection techniques are signature-based recognition and anomaly detection.

Sometimes, a distinction is made between ID and misuse. The term intrusion is used to describe probing or attacks from outside the local network by an unauthorized user, whereas misuse usually describes an intentional or unintentional disruption of service that originates from an authorized user on the internal network.

Host-based ID focuses on monitoring log files and system activity as data sources or sensors for both PC and workstation platforms. A host-based ID tool may monitor network connections and probes to the machine as well as checking and monitoring the file system, process activity, and user activity.

A properly configured host-based ID system will not only secure the user’s working environment from external sources of attacks but will help defend the organization from an insider attack by alerting analyst to a deviation in the user’s activity. Host-based detection also allows for organizational policy enforcement on users by monitoring their activity for unauthorized access to system resources or by logging suspicious data transmissions to suspect IP locations. However, any recordable media will provide an avenue for information theft that cannot be detected by today’s IA tools. In short, a host-based ID system is only one element of an enterprise ID system.

In contrast, a network-based ID system monitors the traffic on its network segment as a data source. A computing platform that is dedicated to capturing IP packets is the data generator or sensor. The incoming packets are compared against a set of predefined data strings or *signatures*. When an incoming data stream matches a signature, an event is generated that must be addressed by human intervention. The ability of a network-based

ID system to detect attacks rests with the analyst's skill in "tuning" the configuration of the ID system with regards to its operating environment and in maintaining an updated database of attack signatures.

Another ID tool working to prevent unauthorized access into the local network from the Internet is the firewall. A firewall is a system that is configured to enforce an access control policy between the local network and the Internet. Firewalls are generally configured to block access by known attackers, or "bad IPs," and to block requests from outsiders for services that are known to be susceptible to attacks. Again, it is necessary for an analyst to configure and maintain the firewall to block new attackers and attack scripts.

Two additional areas that must be mentioned are virus protection and site policies. Both are straightforward, fairly simple to implement, mature, well documented, and should be completed prior to establishing an Internet connection.

Virus protection software has become a critical element in the survival of every computing system with an Internet connection. At the current rate at which new e-mail virus attacks are generated and spread, no system running an Internet connection program could hope to survive without a virus protection program that is frequently updated.

The site policy is an evolving document that defines how and what resources are to be supported in the enterprise and who is responsible for their use and accountable for their misuse. The site policy also establishes the acceptable and non-acceptable practices for users and identifies the security practices that will be enforced by each of the three monitoring technologies, host-based, network-based, or firewall, deployed in the enterprise.

An organization will deploy a mixture of host-based ID, network-based ID, and at least one firewall within its enterprise with an end goal of implementing an automatic monitoring and alerting capability that implements as much of the site policy as possible. While this mixture seems to provide total coverage of the enterprise, there are still many gaps. For example, the sheer numbers of events detected and reported by the tools may overwhelm an analyst and actually provide cover for sophisticated attacks upon the enterprise. This should not be surprising since many ID tools have a high false/positive rate.

Also, while the ID tools appear to be coordinated and functioning as a series of boundaries working together to monitor computer events and network traffic for suspicious activity, in reality they are operating independently of each other. Each tool has been designed to monitor for a specific type of activity and does not take into account what events may have been detected by another component of the ID system. Hence, the data fusion and correlation function is a manual process that rests on the analyst's

shoulders. A disconnected system such as this leaves little room to question the lack of precision and timeliness in detecting intrusion events.

The only thread of a connection among the technologies are the configuration settings made by the analyst, which use his ability to filter out bad events from a flood of harmless activity detected by the ID tools. His knowledge of the enterprise and its usage are captured and implemented in the configuration settings files in the appropriate component of the ID system or in the site policy for the users to adhere to. Even with a finely tuned ID system that eliminates many false alarms, an analyst must manually study the report logs from three different types of tools and possibly the reports from many individual hosts within the enterprise.

Even after the analyst has tuned his ID system and developed tools to help analyze the reported data, he is still at step one. Every time an OS is updated, a new application loaded, or additional hosts added, he must revisit his configuration files to ensure that they provide ID coverage for any new vulnerabilities that may have been introduced as a result of changes to the enterprise. Of course, he doesn't know what those vulnerabilities are until he or someone else has been compromised and reports of the vulnerability are transmitted to the community.

To summarize, today's ID system is a formalized plan of which the data collection components are automated. The man-hour-intense process of configuration, maintenance, data analysis, and correlation are required for daily operation of the system. To date, no COTS automated tools exist to support the analyst in his job of uncovering threats to the enterprise.

3.3.2 Knowledge Extraction

ID technologies provide large volumes of raw data that must be enhanced and analyzed. The field of knowledge discovery and data mining (KDD) provides a rich set of tools for extracting useful information from large volumes of data [FaPSSm96, Grot01, GoJe98]. KDD blends research in various fields such as databases, machine learning, pattern recognition, statistics, artificial intelligence, reasoning with uncertainty, knowledge acquisition for expert systems, data visualization, machine discovery, scientific discovery, information retrieval, and high-performance computing. KDD is sometimes called "Knowledge Discovery in Databases." Its research results and tools are reviewed in this section.

KDD for ID has been researched and implemented in the fields of radar intrusion [NeFi96], telecommunications [Copr01, FaPr97, Grot01], marketing [HKMY98, Grot01], and medicine [BTTh99, Grot01]. User profiling, which aims to collect useful information pertaining to each individual user, has been researched in e-commerce Web-based applications and other marketing applications [KoPr01]. User profiling refers to

constructing accurate and comprehensive profiles that describe important information such as who the customers are and how they behave [AdTu01]. Each of these fields has contributed certain techniques that could be applied to computer intrusion. In particular, user-profiling techniques can be used to construct a profile of each intruder.

The KDD process consists of several steps: data selection, data preparation, data cleaning, incorporation of prior knowledge, data mining, and proper interpretation of the mining results. The main component of the KDD process is data mining.

Data mining is defined as the process of extracting descriptive models from large stores of data. It involves fitting models to data or extracting patterns from observed data. Data mining algorithms consist largely of a particular mix of three components: the model, the selection (i.e., preference) criterion, and the search algorithm.

Major approaches for data mining include [FaPSSh99, LeStMo98, MePs98]

- *Classification* of data into predetermined categories (i.e., “bins”).
- *Link analysis*, which determines relations (such as association rules) between fields in database records.
- *Sequence analysis*, which models sequential patterns and time-based sequences of events. A sequential pattern is an association between sets of items, in which some temporal properties between items in each set and between sets are satisfied. Items in a set have the same temporal reference, and an order between sets is established by means of the temporal reference.
- *Similarities* in ordered data, such as clustering or dependency modeling.
- Summarization.

The most popular data mining models include statistical modeling such as Chi-square tests and regression analysis, decision trees, rules, linear models, non-linear models (such as neural networks and genetic algorithms), example-based methods (such as case-based reasoning), probabilistic dependency models (such as Bayesian networks), and relational attribute models. A review of various techniques is provided in Section 3.3.4. Search algorithms are of two types: parameter search given a model and model search over model space. SAS SEMMA method is used in most data mining software: *Sample* the data, *Explore* the data, *Modify* the data, *Model* the data, *Assess* the data.

3.3.2.1 User Profiling and Association Rules

One of the most useful concepts of data mining that pertains to C-IA decision-making is user profiling. There are two major approaches to user profiling. In the first approach, profiles are constructed from customer’s demographic and transactional data and contain factual information. In the second approach, customer profiles contain factual data as well as behavioral rules. Rules can be either specified by domain experts, in which case they apply not to individual customers but to groups of customers, or derived from user

transactional data using data mining methods [AdTu01]. For example, user profiling can detect patterns in users' Web browsing.

A commonly used approach to describe behavioral patterns is by using *association rules*. A sample association rule might be

"if this consumer buys item X on weekends, he usually buys
item Y as well."

In the field of computer security, a sample association rule might be

"if the attacker X attacks from IP address Y,
he usually attacks after 9pm and before 12am on port Z."

The aim of the association rules is to extract behavioral information from data. One of the most relevant problems in data mining is the discovery of association rules by mining the collected data.

An association rule has the form $X \Rightarrow Y$, where X and Y are two sets of items. Data mining usually produces several association rules, so less useful rules should be eliminated. Two common metrics used to measure the "goodness" of individual rules are *support* and *confidence*. Therefore, an association rule is usually given in the form

$(X \Rightarrow Y, \text{confidence}, \text{support})$.

If N is the total number of samples in a data set, then for a rule of the form $X \Rightarrow Y$, we define as follows:

- *Support* measures how much of the data set the rule covers.
- $\text{Support} = P(Y) = (\text{number of occurrences of } Y)/N$
- *Confidence* measures the correlation between the antecedent and the consequent of the rule. It is the conditional probability to find Y in a group given X was found in the group.
- $\text{Confidence} = P(X \text{ and } Y) / P(Y)$.

For example, an association rule for a customer might be

$(\text{buy newspaper} \Rightarrow \text{buy coffee}, 0.4, 0.1)$.

This rule indicates that this user buys coffee 40% of the time when he buys the newspaper and that buying coffee constitutes 10% of the activities recorded for this user. Separate thresholds for support and confidence are given by the user to discard the less frequent association rules. Further processing of association rules includes discovering recurrent (or frequent) episodes [MaToVe97].

Current research in data mining indicates that using support and confidence with association rules might not be effective, because a rule might be discarded if the calculated support is low. However, the fact that support is low may be significant, if we

expect support to be high. [Grot01] gives the example of Pepsi and Coke purchases. Assume that each Pepsi and Coke purchase occurs with a 50% chance. If actual data confirms that Pepsi and Coke purchases occur together in 1% of purchases although we expected 25%, this data point should not be discarded. Low support can indicate many scenarios, such as Pepsi and Coke are bought by a different set of consumers, or they are bought by the same set of consumers but either one or the other is bought depending on, for example, current promotions. Another question that needs to be answered is whether the events happen together by chance, randomly, or they are truly correlated. [Gort01] suggests using Chi-square tests to determine statistically significant non-random associations.

Once rules and user profiles are determined, the next step is to perform prediction of what the user is most likely to do next or what the user's response might be. Regression techniques can be used.

3.3.2.2 KDD Applications to Intrusion Detection

KDD has been applied to intrusion assurance mostly for detecting fraud in telecommunications, financial, and law enforcement areas. Telecommunications applications are most relevant to C-IA.

The telecommunications industry has developed some ID procedures to detect fraudulent telephone use. Procedures were developed to detect intruders who illicitly access legitimate users' accounts (in the telecommunications industry, this attack is called *cloning*) [FaPr97]. The ID is based on constructing a user profile for each legitimate user by assigning association rules. The solution described in [FaPr97] could be suitable for the C-IA application only to a small degree, because it uses very specific mobile telephone industry metrics, such as the physical distance traveled.

However, fraud detection in telecommunications as described in [CoPr01] is rather relevant to the C-IA application, because it deals with detecting patterns in data streams, i.e., user transactions. If [CoPr01] concepts are applied to the computer intrusion, the following information can be extracted: probing rate, proportion of attacks to specific targets, fuzzy classification of attack duration, fuzzy classification of attack hour and day (work vs. weekend), fuzzy classification of the top countries, or recently attacked IPs. Other useful concepts from this paper include ideas for assigning a general signature to a newly discovered user or attacker and then learning the user's behavior.

Based on the published literature, it appears that the telecommunications industry predominantly uses in-house KDD tools. Other industries, such as financial and law enforcement, seem to mostly rely on commercial packages. Numerous and diverse commercial KDD packages are available [Grot00]. For example, CaseRunner [CaseR] is designed to help investigators build a chain of evidence. The tool searches through data

to build visual connections between data, such as people and events. In other words, CaseRunner tries to visually represent “a story,” for example, “person A owns bank account X, drove to city Y to meet person B, exchanged item C.”

Many companies provide fraud consulting and software, such as SAS [SAS] and HNC [HNC], typically for credit, insurance, financial, and telecom fraud. SAS SEMMA method (see Section 3.3.1) is used in most data mining software. SAS has customers in all business sectors, especially because SAS gives inexpensive licenses to academic institutions. HNC has a product called Falcon [Falcon] that uses neural network to detect fraud. Twenty top-level credit card companies use Falcon to analyze more than 400 million credit cards. HNC offers various software packages for insurance, telecommunications, financial, software management, and government tax-collection applications. Financial markets have used neural networks by IBM, SAS, SPSS, HNC, Angoss, RightPoint, Thinking Machines, and NeoVista.

Some of the most widely used KDD software packages are listed in Table 3 [Grot00].

Table 3 — Widely Used KDD Software

Application	Company	Tool	Methods Used
Data Mining	SAS Institute	Enterprise Miner	Regression, decision trees, neural network
Estimation problems	Script Software	KnowledgeMiner	
Estimation for stocks	Neural Applications	NETROPHNET	
Fraud detection by prediction, financial, and credit card fraud	HNC	Falcon	Neural network
Marketing	Informix Software (Right Point)	Real Time Marketing Suite, Data Cruncher	

A small portion of KDD research has been targeted to ID and response for security applications [ILKe95, LeSt98, LSC01, LSM98, LSM99, LPS99], mostly applying the existing KDD research such as user profiling and association rules.

3.3.3 Situation and Risk Assessment

Once the data about a possible attack presence is obtained through the process of knowledge extraction, knowledge engineering is applied to obtain the “big picture,” assess the situation, and provide risk analysis.

Ideally, the situation assessment (SA) module provides all of the following information [CSAP21]:

- *Attack profile*, the method (such as techniques, tools, utilities, and steps taken)
- Intruder sophistication
- *Readiness* to respond to the threat
- *Potential courses of action* (CoAs), with description and rationale

The SA module's output can include attack profile, INFOCON, safeguard options, system change, system statistics, and network update options such as recommended network tools, equipment, and configuration, along with automatic response capabilities.

Ideally, the risk analysis module (RA) provides all the following outputs [CSAP21]:

- Best CoA
- Potential risks to the target system and related mission
- Prediction of the next target
- Threat profile and risk

In the state of the art in current practice, the manual top-down approach to situation and risk analysis involves manual assessment:

- Manual examination of ID tools' output (i.e. sensor alerts) to determine the most critical alarms.
- Manual examination of the most critical alerts to determine patterns of misuse, taking into consideration various "soft" variables such as the current political situation.

The cognitive steps used by security analysts in the field are outlined in the subsection below.

Once the situation assessment and risk analysis are performed, remedial actions are executed, usually manually. Alerts to other geographically distributed networks are also manually executed, for example, via a phone call.

3.3.3.1 Threat Assessment

The following procedures are cited from [Bora01], specifically, the section "Top down approach on improving security."

"The situation assessment and risk analysis steps used today by security analysts in the field include

- Asset analysis. What needs to be protected? Information and processes are listed, for example, What are the important assets? Are they stored on computer? What are the cost implications of loss of these assets? The measures taken to protect assets should correspond to the value of assets.

- Analysis of current security rules/policies/practices (if any).
- Defining basic security objectives, such as basic availability, confidentiality, and integrity objectives.
- Threat analysis. Before deciding how to protect a system, it is necessary to know what the system is to be protected against i.e. what threats are to be countered. Threats are identified (employee vengeance, hackers, espionage, technical failures etc.). A list of sample threats is presented in the Appendix.
- Impact analysis. The impact should be judged by decision-makers, not technical experts.
 - What is the impact or consequence (harm to organization) if a threat, or a combination of threats, is realized? The impact is specific to the organization, for example, loss of company secrets, modification of accounting data, falsification of money transfers.
 - The impact has two components, a short-term impact (threat is short) and a long-term impact (the threat persists, affecting the business in the long term). The total impact should be considered as a number (0-5) with a contribution for the short term and the long term.
 - The impact is negligible.
 - The effect is minor; major mission operations are not affected.
 - Mission operations are unavailable for a certain amount of time, revenue is lost, user confidence is affected minimally (unlikely to lose customers).
 - Significant loss to business operations or customer confidence or market share; customers will be lost.
 - The effect is disastrous, but the company can survive, at a significant cost.
 - The effect is catastrophic, the company cannot survive.
- Calculate risk.
 - What is the *likelihood* of a threat occurring (0-5)? Technical experts can probably judge better than business experts what the likelihood of a threat occurring is:
 - The threat is highly unlikely to occur.
 - The threat is likely to occur less than once per year.
 - The threat is likely to occur once per year.
 - The threat is likely to occur once per month.
 - The event is likely to occur once per week.
 - The event is likely to occur daily.
- $\text{risk} = \text{impact} * \text{likelihood}$
 - The risk can have a minimum value of 0 (no risk) and a maximum of 25 (extremely dangerous risk). The greater the risk value, the more important it is to implement counter measures.
 - An acceptable risk value needs to be set. All risks having a value higher than this number are unacceptable risks that must be countered. For this project, we (provisionally) set the acceptable risk = 15.

- Constraints analysis: Examine requirements outside of your control (national and international laws, corporate requirements, corporate culture, contractual requirements, budget).
 - Decide on a counterstrategy.
 - Define security objectives.
 - Define countermeasures (e.g. policy, roles, processes, responsibility, mechanisms).
 - Can risks be reduced to an acceptable level with this strategy? Are costs too high?
 - If not, can the remaining risk be economically insured?
 - Otherwise, redo the strategy.
- Implementation:
 - Develop security policy and guidelines together with an information classification system.
 - Define a security organization (or modify the current organization). Users, administrators, and managers should have clearly defined roles/responsibilities and be aware of them.
 - Run pilot tests. Tune policies, processes, and organization according to results.
 - Secure systems on a wide scale.
- Assurance: Reevaluate risks and security strategy regularly.”

According to [Boro01], threats are divided into the following categories: General, Identification/Authentication, Availability, Privacy, Integrity/Accuracy, Access Control, Repudiation, and Legal. For each type of threat, a table is presented showing the threat description, the impact of the threat (with the score 0-5), and the likelihood of the threat occurring (with the score 0-5). The reader is highly encouraged to read these tables in [Boro01], which present an excellent overview of the threat-assessment questions asked by security analysts in the field.

3.3.4 Expert and Knowledge-based Systems

Expert and knowledge-based systems are used to make decisions, and thus present the available foundation on which to build C-IA policies and decision-making modules, such as information gathering, correlation, and decision-making and coordination performed on and between OversightNets, OverNets, and OuterNet.

Expert and knowledge-based systems are usually based on a set of rules and implemented on the basis of one or more of the following:

- Symbolic method
- Neural network (also called connectionist model)
- Fuzzy sets
- Other

Most systems are hybrid. For example, we can use:

- A neural network or genetic algorithm to extract and/or refine fuzzy rules
- Fuzzy rules or genetic algorithms to optimize the weights of neural networks and/or determine neural network structure
- Neural networks and fuzzy rules in symbolic systems

3.3.4.1 Symbolic Methods

Symbolic methods use traditional, deterministic mathematics such as differential equations, propositional logic, and predicate logic. Symbolic systems based on logic can use IF-THEN rules. Connective operators allowed are OR, AND, NOT, $\rightarrow$, and $=$. Propositional logic does not allow the use of variables.

Symbolic methods cannot deal with uncertainty. For example, logic cannot be used to conclude that `Jim is mortal` given the rule and the fact

```
IF human THEN mortal.
```

```
Jim is human.
```

However, this kind of reasoning is used in expert systems for applications such as diagnosing. Therefore, we use probabilities, certainty factors, and possibility/necessity concepts to deal with uncertainty.

3.3.4.1.1 Rules

Rules are in the form

```
IF condition THEN conclusion.
```

For example,

```
IF (temperature > 38°C) THEN (take aspirin).
```

(A rule in the above form, where the subject is transformed, is called a production rule.)

3.3.4.1.2 Decision Trees

Decision trees partition the problem space according to a set of criteria and “guide” you to the solution. The final decision is based on partitioning criteria, which may lead to the wrong solution. For example, if patients are characterized by regularity of heart rate and blood pressure, the decision tree can partition on heart rate and then on pressure, or vice versa.




Figure 1 — Decision Tree Example

The paths in a decision tree can be represented as rules. A technique called inductive decision tree can build a tree based on a set of training data.

Rules and decision trees are quite similar, except that rules can be used to describe relationships between variables in general. For example, a rule might be

IF (source IP = destination IP) THEN ...

where it is not necessary to specify exact values of source IP and destination IP. In decision trees, the exact values have to be specified.

3.3.4.2 Neural Networks (Connectionist Model)

Neural networks are networks of “neurons” that fire on a given input; inputs have “weights,” which can be programmed in or learned based on a set of existing (i.e. training) data.

Good for machine learning, generalization (can approximate any function), massive parallelism (neurons work in parallel), robustness (even some neurons make mistakes, the overall solution will not suffer), and partial match.

Neural networks do not remember the reasoning process. They can work poorly if the data set is large. They can also be over-trained, i.e., they work well only on the training data set.

3.3.4.3 Fuzzy Systems

A fuzzy system consists of

- fuzzy input and output variables

- a set of fuzzy rules
- fuzzy inference mechanisms

Fuzzy variables have “fuzzy,” or descriptive, values that overlap. For example, it is possible to be both “young” and “old” to a certain degree, called a “truth degree.” Each fuzzy variable has a truth function μ associated with it.

For example, a fuzzy system has the following rules:

IF (a person is young) THEN (they eat a lot of pizza).

IF (a person is old) THEN (they eat some pizza).

Since a fuzzy variable can have several values, several rules can fire at the same time, and the conclusion is obtained using an inference mechanism.

3.3.4.4 Case-based Reasoning

Case-based reasoning relies on a past solution to determine a present solution. Solutions to old problems are stored, and the most suitable solution is selected and adapted to the new problem. Case-based reasoning is most suitable for legal and other dispute/mediation-oriented applications. Other suitable application domains include medicine, cooking, and process control and engineering design.

3.3.4.5 Stochastic Search Methods

3.3.4.5.1 Non-evolutionary

In this category, we put different types of stochastic search methods, such as gradient method or simulated annealing.

3.3.4.5.2 Genetic Algorithms

Evolutionary search. The algorithm is as follows:

Initialize population of possible solutions

While a criterion for termination is not reached,

 Crossover two specimens

 Potentially change the resulting specimens

 Select the most promising ones

Good for search and optimization problems, such as:

- Optimizing parameters
- Optimizing neural network architecture and parameters
- Optimizing fuzzy rules

3.3.4.5.3 Evolutionary Strategies

Evolutionary strategies involve the same principle as genetic algorithms, but change and selection of specimens is based on statistical and other characteristics, not binary vector.

3.3.4.6 Statistical Methods

Some simple statistical properties are mean (expected value), standard deviation, variance, correlation, and covariance.

Mean: the value that all data points tend to.

Standard deviation: measure of how much data deviates from the mean.

Correlation: measure of how much two variables depend on each other.

3.3.4.6.1 Generalized Linear Models (Regression)

Determines how one variable y is related to other variables $x_1, \dots, x_N$. The most widely used form of regression model is general linear model:

$$y_j = A_0 + A_1 x_{1j} + \dots + A_n x_{nj} + e_j, \quad j = 1, \dots, m.$$

e_j must be independent and distributed as normalized Gaussian.

Can use least square estimators for A 's maximum likelihood estimators. Can use ANOVA to estimate which A_i are non-zero.

Regression can be performed on non-linear models, or generalized linear models.

3.3.4.6.2 Chi-square Test

Used to determine how much observed data (o) deviates from value we expected (e).

$$\chi^2 = \sum_{i=1, N} (o_i - e)^2 / e$$

Using the calculated χ^2 statistics and Chi-square probability distribution, it is determined if the deviation from observed data is based on chance (i.e. the observed data is behaving as expected) or not (i.e. the observed data does not match the initial hypothesis.)

3.3.4.6.3 K-nearest Neighbor

Based on evaluating distance between objects, such as the following:

$$\text{Absolute distance } D = \sum_{i=1, N} |a_i - b_i|$$

$$\text{Euclidian distance } E = \text{sqrt}(\sum_{i=1, N} (a_i - b_i)^2)$$

3.3.4.7 Dealing with Uncertainty

3.3.4.7.1 Confidence Factors

Express belief that a fact holds true. CF of -1 indicates complete disbelief, CF of $+1$ indicates complete belief. CF is attached to the rule conclusion, indicating how much belief we have that this rule holds. For example,

```
IF    (current economic situation is good
      AND market is going up)
      THEN buy.  CF=0.9.
```

Other parameters can be attached to the rule conclusion: degree of importance, sensitivity, noise tolerance, etc.

3.3.4.7.2 Possibility/Necessity, Similarity

Possibility is the degree to which an expert considers a hypothesis H to be feasible or possible. Possibility is different than probability because it is a non-statistical concept, which represents capacity or capability. For example, the possibility of throwing a die and getting 6 is 1, but the probability is $1/6$. The possibility of throwing a die and not getting a 6 is also 1.

Similarity S is the measure of how much fuzzy set B matches fuzzy set A .

$$S = P(A/B) \text{ if } N(A/B) > 0.5$$

$$S = (N(A/B) + 0.5) * P(A/B), \text{ otherwise}$$

where P is the possibility and N is the necessity.

$$P(A/B) = \max(\min(m_A(x), m_B(x))) \text{ for all } x.$$

$$N(A/B) = P(\text{not } A/B)$$

where m_A and m_B are truth functions.

3.3.4.7.3 Probabilistic Methods

Markov chains, Bayesian networks, and simulation are examples of probabilistic methods, where they all have the same weakness, which is favoring the most possible outcome.

Bayesian Belief Networks (BBNs)

Based on conditional (Bayesian) probabilities. Bayes' theorem states that probability of X happening given that Y happened is

$$P(X/Y) = P(X \text{ and } Y) / P(Y)$$

A BBN is graphically represented as a network of variables, called a Directed Acyclic Graph (DAG), which defines a model of conditional dependencies among the variables. BBN can use statistical methods to learn the conditional dependencies from training data.

3.3.4.8 Approaches at a Glance

Table 4 — Approaches at a Glance

Method	Suitable Use	Disadvantages	Notes
Statistical	Statistically representable data is available, and the underlying type of goal function is known.	Very often requires simplification, such as assuming that system is linear or of some known form.	
Symbolic AI rule-based system	Problem knowledge is in the form of well-defined, rigid rules.	Adaptation of rules is either impossible or very difficult.	Good when the problem is rich in theory and poor in data.
Fuzzy system	Problem knowledge includes heuristic rules that are vague, ill-defined, approximate, possibly contradictory.		
Neural networks	Problem knowledge includes data without having any knowledge as to what type the goal function might be. Can be used to learn heuristic rules after training with data. Can be used to implement existing fuzzy or symbolic rules.	Unless enhanced, neural net does not remember how it came to the solution. Difficult, if not impossible, to represent variables.	
Genetic	Very efficient when only little	Input must be a bit-	

Method	Suitable Use	Disadvantages	Notes
algorithms	is known to start with. Requires neither data sets nor heuristic rules, but only a simple selection criterion to start with.	vector.	




Figure 2 – Different Pathways to the Solution, as suggested by [Kasa98], p. 66.

3.4 C-IA Architecture

The C-IAA, which was conceived prior to the start of this project, was reviewed and analyzed for feasibility from two perspectives: as an architecture that envisions a system and as a collection of abstract components that envision their respective hardware and software implementations. This section describes the C-IAA and its constituent abstract components from the systemic viewpoint. Implementation issues and recommended research are identified at the system level. Later sections address these considerations for the various abstract components that constitute the C-IAA.

3.4.1 Concept

The C-IAA partitions the concurrent IA problem into a series of hierarchical domains. Each domain has a set of IA responsibilities that is appropriate to its span of control. An abstract block diagram that depicts the elements of C-IAA is provided in Figure 3.




Figure 3 — C-IAA Abstract Block Diagram

ApplicationNet. Beginning at the lower left of the diagram is the purpose of all IA, the applications for which the processing infrastructure exists. C-IAA makes few assumptions about the applications it supports, except with regard to structural flexibility. Applications may be centralized or distributed, monolithic or otherwise, physically collocated or remote. The communications services and networks that the applications use to perform their business is called the ApplicationNet in C-IAA. It is the ApplicationNet that is assumed to be subject to attack outside or inside the organizational enclave. The ApplicationNet is an abstract element that may consist of one or more physical networks.

Processing nodes. In C-IAA, IA security begins on the same platforms as the applications to duly monitor and control the application environment. C-IAA asserts that it is impossible to provide host-based IA, if the responsible IA components are subject to the same attacks as the applications. Therefore, C-IAA postulates that IA functions must be segregated from the applications within a secure computing environment. Once securely isolated, IA components can perform meaningful host-based ID, Analysis, and Response (IDAR) on behalf of vulnerable applications. Traditional network-based IDAR, which is usually performed using hardened platforms that do not support general applications processing, is also an important ingredient of C-IAA, but has been omitted from the diagram for the sake of clarity.

OversightNet. At the second level of the hierarchy, C-IAA proposes that the IA activities performed on the individual processing nodes can be coordinated to yield a concurrent and centralized IA for the organizational enclave. Whether centralized on a special-purpose platform or distributed among the IA components of the applications nodes, this coordinated IA would enable the domain of the organization to detect more attacks and with greater certainty than is possible with standalone detection. Moreover, the greater data and centralized viewpoint provides IA managers of the domain with the opportunity to mount a coordinated response to a diffuse attack. At this level, C-IAA postulates the introduction of automated attack detection, correlation, assessment, and response tools. To further isolate the IA security components from possible attacks, communications among the components is segregated using services and networks called the OversightNet. Like the ApplicationNet, the OversightNet is an abstract concept, and implies only that IA security components can communicate securely with privacy, integrity, and authentication of correspondents. The physical transport that underlies the OversightNet could be physically separate from the ApplicationNet or the same transport on which high-quality security services have been applied. Occasionally, the term OversightNet is used to mean the collection of coordinated C-IA functions at the organizational level rather than the communications between IA security components. Care should be taken to maintain the distinction between secure communications and the coordination of C-IA functions.

OverNet. In the real world, organizations can be relatively flat with centralized responsibilities for applications support or quite hierarchical with delegated areas of responsibilities and associated resources. Both types of organizations are increasingly required to interface with other organizations with similar interests to achieve a common goal. Whether between the hierarchies of a single complex organization or between different cooperating organizations or some combination thereof, C-IAA postulates the ability to coordinate the IA functions and services between the entities using the concept of the OverNet. In the same way that the OversightNet centralizes the command and control of an organization's IA actions, the OverNet centralizes the IA actions of multiple organizations. At this level, a more sophisticated form of attack detection, correlation, threat assessment, and response is anticipated. However, greater challenges lie in maintaining organizational integrity. The OverNet also presupposes a means of secure communications among the OversightNet functions of the participating organizations because physical remoteness is likely. In addition, the OverNet must contend with mobile participants as well as participants that join and leave the OverNet as situations change.

OuterNet. C-IAA contemplates the last and highest level of the hierarchy with three different IA support functions in mind: 1) coordination of IA activities among multiple OverNets (thereby enabling the recognition of a coordinated attack against two different industries, for example), 2) standardization of emergency response activities when an OverNet has been violated, and 3) the widespread monitoring of public-access networks with ID and analysis tools. Because the OuterNet would itself become a major target of attack, it has been anticipated that communications support for the OuterNet would be highly or totally isolated from all low-assurance networks, public, private, or governmental.

Sample applications of the C-IAA might be

- Combat situations in which weapons and troops are synchronized during attack, such as in DARPA's Networked Targeting Technology program [NTT00]. Such situations can involve manned and unmanned operations and real-time, near-real-time planning, and re-planning.
- Defense against information attack such as uncoordinated and coordinated attacks by viruses, worms, denial of service, gaining of administrative privileges, network probing, and various other types of attacks.

3.4.2 C-IA Requirements

In many ways, C-IAA is an amalgam of concepts and technologies that have appeared in other contexts. Yet, the concept of C-IA itself imposes a series of systemic challenges that set the C-IAA apart. At its heart, C-IAA implies, in combination, the following requirements in its constituent components:

- High-assurance segregation of trusted security functions

- Cooperative engagement among multiple asynchronous functions
- Optional participation by entities at all hierarchical levels
- Advanced, automated tools for the analysis of attacks, threats, and responses
- Methods for negotiating and coordinating policies among entities
- Common methods of expression for relevant data, policies, commands, and status
- Methods of maintaining organizational independence, integrity, and privacy while achieving meaningful coordination

3.4.3 Implementation Issues

Implementation issues for the architecture as a whole are, for the most part, the sum of the implementation issues of its constituent abstract components, each of which is addressed in sections following. However, some issues transcend the individual components. These include

- Consistent security policies for IA, from which security functional, isolation, and packaging requirements can be derived.
- Security policies governing the exchange of IA data among separate organizations.
- Research and development tools that support the development of secure and trusted software components.
- Guidelines for the incorporation of COTS products into the C-IAA that address risk–benefit analysis.
- Guidelines for the development of performance objectives, including such factors as events-per-unit-time, correlation capacity, decision response times, CoA dissemination and response times, and certainty objectives.

3.4.4 Recommended Research

The following areas of recommended research address the C-IAA as a whole and are above and beyond the more specific recommendations that address the needs of the individual abstract components presented elsewhere.

3.4.4.1 Security Policies Governing Information Assurance

Using the work surrounding the protection of data for national security and other applications such as banking and personal privacy, a project should undertake the development of prototype security policies for IA that would enable the multiorganization, hierarchical C-IAA, as currently envisioned, to be deployed as well as implemented.

Using existing infosec and comsec guidelines as templates, the project would attempt to answer such questions as

- When, if ever, does IA data become a matter of national security?
- What portions of IA data can be considered tactical and transient versus strategic and long term?
- What levels of protection should be afforded to IA data as compared to national security data?
- What safeguards, if any, must be applied to IA data when it leaves an organizational enclave?
- Is IA data subject to the constraints of aggregation as the different levels of hierarchy collect them?

The project should strive to develop written prototype policies and guidelines that could be evaluated by the various governing agencies.

3.4.4.2 Secure Software R&D Tools

Software is easy to write and difficult to test. Some existing software is notorious for security holes, for example, common gateway interface (CGI) and server API modules. Guidelines for writing, testing, and evaluating trusted software in the context of C-IAA should be developed prior to implementation.

Based on current trends and the widespread use of such development environments as Java, ActiveX, and C++, a basic research problem is to deal with the difference between executable versus active content. Programs and data are treated equally and stored in the same memory. This policy allows a rogue program, such as a virus, plug-in, or Java applet, to modify both data and programs, eventually modifying itself. Research needs to address how such inappropriate modifications can be prevented or how the resulting damage can be prevented or mitigated.

It remains to be investigated how languages such as JavaScript, Java, and many others can be made more secure and/or used more securely. Some programming languages are less vulnerable than others; for example, Java is known to behave well if given misconfigured input, while C and C++ do not [Opp100]. Many programming languages allow programs to open network connections and thus increase the system's exposure to vulnerabilities. An interpreter for a programming language, for example, Tcl/Tk, can open, modify, or delete any file on a computer system. Other research areas include

- Developing policies for the use of various programming languages and other executable code.
- Enhancing existing programming languages, such as Java and JavaScript, which are still evolving and receptive to the addition of security features.
- Developing trusted libraries.

3.4.4.3 Guidelines for C-IAA Systems Engineering

Unlike mature architectures, C-IAA lacks meaningful parameters or “rules of thumb” that systems engineering efforts can use to guide the bounding of a conforming system or its performance. Research into the operational needs of IA administrators and analysts, their corresponding systems administrators, and the business or command structure they support is necessary to determine such diverse system characteristics as

- The rough order of magnitudes of system response times necessary for usefulness.
- Which processes must be automated before all others to make the system tractable.
- What existing management or operations support systems must be accommodated.
- What interfaces and reports will best expedite the IA response and control process.

3.4.4.4 Evaluating Secure Communications Protocols

More investigation is needed to perform security analysis on existing protocols and software. For example, when security analysis was performed on the SSL protocol, which is widely used for e-commerce, it was found to be safe with minor modifications [WaSc96]. However, the Bleichenbacher and similar attacks based on inherent protocol flaws need to be discovered. Formal methods can be used for top-down risk analysis and implementation of security in an organization. Attacks based on faulty implementation, such as buffer overflows, need to be discovered by red teams.

The objective of this recommended research project would be to identify those secure transport protocols of most use to C-IAA, determine the extent they need to be evaluated and qualified prior to admission into the architecture, and estimate the risk associated with not performing the needed evaluation. The proposed project should also prepare uniform guidelines for the evaluation of transport protocols, including which risks are best exposed by what techniques.

3.5 C-IAA Processing Node

3.5.1 Concept

In C-IAA, the application-processing node is a platform for mission and support functions, which are the reason for the node's existence. As such, it is an obvious target for information warfare attacks. To monitor and constrain such attacks, C-IAA asserts that each such platform must include IA functions in addition to the primary application functions. For the IA functions to be effective, they must be trustworthy and securely isolated from the application environment of the processing node so that they will not be subject to attack, corruption, or subversion. Toward that end, C-IAA has postulated an OversightNet security kernel, which would manage the hardware platform to securely segregate the application and security environments.

Secure methods for software to share the same hardware environment have been under active consideration for over 30 years, with stops, starts, and restarts. As hardware capabilities have increased, the need to share hardware has waxed and waned, and the operating systems that dominate the market have come and gone. As a result, many approaches to secure sharing have been theorized, designed, built, tested, and frequently discarded.

3.5.2 C-IA Requirements

The C-IAA processing node must be capable of

- Supporting existing prime mission applications
- Securely segregating applications processing from security processing
- Executing IA security functions in real time with respect to the frequency of application transactions and IW attacks
- Securely communicating with C-IAA security entities

Of these requirements, this study has focused primarily on the requirement to segregate the applications environment from the security-processing environment. As discussed previously, the need for secure communications is universal within the C-IAA, and many mature, off-the-shelf, products are available to support them. Similarly, the concurrent computing requirement can also be met with a wide variety approaches including multitasking on high-speed monoliths, multiprocessor arrays, and multiprocessor networks, all of which are represented by mature, off-the-shelf products. The requirement to support existing prime mission applications must be deferred until such applications can be identified.

3.5.3 Current State of the Art

Today's market, dominated by Microsoft Windows, Unix, and Linux, offers three basic approaches to secure hardware sharing:

- Secure operating systems, which are generally built on secure kernel technology and can be extended to create a Trusted Computing Base (TCB),
- Security patches to non-secure operating systems, and
- Trusted co-processors, which feature a combination of a TCB with varying degrees of physical hardware protection.

Each of these approaches segregates sensitive processes and data so that a process must possess sufficient express authorization to access or modify them. All the classic techniques are represented, including access control lists, labeling, and implementations of mandatory and discretionary access.

The challenge posed by C-IAA is selecting an approach for the processing node that is sufficiently robust while meeting the other three processing node requirements.

3.5.3.1 Approaches Related to Kernels

Approaches for achieving a TCB with existing kernels include the following:

- Patches on top of existing kernel, i.e. kernel hypervisors, which load on top of an existing kernel and monitor to ensure that a policy is not violated. These patches are installed separately from the kernel and application software and do not require any changes in the kernel or software. This is often the most practical solution. However, if the kernel or the application software change, the patches must be updated. In the simplest case, patches are not aware of each other, and the user cannot configure them. A patch is usually intended to work with a specific application.
- Running copies of the existing secure kernel for each class of user privileges or running a multilevel operating system using a proxy time configuration.
- Designing a new secure kernel.
- Using trusted architectures.

The Linux community has developed many security patches for the Linux kernel, including the following:

- Linux ID System patch [LIDS]. LIDS is a patch and a set of administrative tools that implements a reference monitor and mandatory access control in the Linux kernel. It prevents unauthorized execution and process destruction, as applied to file systems, directories, hard disk, and RAW IO. The idea behind LIDS is that even the super user (the root) can be prevented from executing unauthorized actions. LIDS logs all unauthorized actions and can shutdown the unauthorized user's session at once.

- Rule-based access control [LDP].
- Low watermark mandatory access control [LDP].
- Audit daemon [LDP].
- Security patch to prevent buffer overflow, restricting creation of hard links, etc. [Openwall, LDP, UFN].
- Access control lists and mandatory access control [Privs]. Linux-Privs is based on the POSIX.1e security model.

A more complex kernel hypervisor approach was developed [MiLuBr97] in which a client kernel hypervisor interacts with various applications, a master kernel hypervisor controls client kernel hypervisors, and the client hypervisor module allows the user to configure client hypervisors.

Several secure kernels and operating systems exist:

- Trustix Secure Linux 1.5 (TSL) [Trustix] is a secure Linux OS specifically written for Linux servers. TSL is intended to provide maximum security with a minimal number of services and minimal installation requirements. For example, TLS does not support X-windows. TLS is publicly downloadable.
- Openwall GNU/*/Linux (Owl) is a security-enhanced operating system intended for servers. It is based on Linux and GNU as its core, compatible with other major distributions of GNU/*/Linux. It is publicly available [Owl].
- sLinux [sLinux] is a secure version of the Linux kernel. sLinux is intended to be a secure, specialized, server distribution of Linux.
- Security-Enhanced Linux (SELinux) is a secure Linux OS by NSA. NSA, SCC, and the University of Utah produced a Flask secure operating system, which is being incorporated into the Linux operating system to produce SELinux. NSA recently partnered with PGP Security (a division of Network Associates) to continue working on SELinux. SELinux includes access-control mechanisms in the kernel that helps prevent security breaches at the application level. Since “federal policy has called for increasing the federal government’s role as both a user and contributor of open-source software [Bitta01],” SELinux is available to the public under the general GNU public license terms [SELinux].
- Type Enforcement is a trusted operating system that provides mandatory access control, confinement, and least privilege. NSA contracted SCC to develop a “robust, secure version of Linux” [Penn00]. SCC has developed proprietary “Type Enforcement Technology,” which operates at the lowest level of Linux kernel. Type Enforcement was used for SCC’s Unix firewall product, called Sidewinder. Sidewinder has not been hacked after 23,239 attacks, as of July 2001 [Sidewn].
- Advanced Infosec Modules (AIM) technology by Motorola is based on the separating kernel concept, i.e., a kernel that strictly separates processes and

allows different algorithms to run simultaneously without unwanted interaction among them. Thus, separation kernels allow for running multiple levels of security using programmable cryptography. AIM technology incorporates a range of programmable capabilities in a single processing chip, called an AIM chip, and relies on MASK functionality to provide the secure operating system and guarantee separation of all data for multilevel security. The NSA certification process certifies the core cryptographic capability. Additional modules can be programmed in by the customers and certified in less time. The AIM separation kernel was formally specified and verified using formal methods and a tool called Specware [MWV00].

- Trusted Computer Solutions, Inc. has developed a multilevel operating system (Trusted Solaris/HPUX-CMW) using a proxy type configuration. The proxy server is a centralized point of access for internal top secret/secret networks and runs on a compartmentalized workstation. [Perry98].

The Open Software Foundation (OSF) GNU HURD [HURD] is a free substitute for the Unix kernel. HURD is object-oriented and extensible. It remains to be investigated how it can be extended into a security kernel.

Arizona University was funded by ARPA until 1996 for the Highly Structured Architecture for Network Security (HOSANA) [HOSANA]. HOSANA's research is applicable to the ApplicationNet/OversightNet concept. HOSANA developed an x-kernel environment as a Linux IPSec environment. Each HOSANA host has an application security management module for managing what C-IAA terms the ApplicationNet and a crypto-enhancement module and security management module for managing what C-IAA terms the OversightNet.

3.5.3.2 Approaches not Related to Kernels

Compartmented Mode Workstation (CMW) security is an architecture in which a trusted client, typically a window manager, controls applications and processes. This client maintains labels containing permission levels for all windows, X atoms, and pieces of data for all clients and servers. These labels are checked for each transaction, and a detailed audit trail is maintained. CMW implementations have been certified as meeting the NSA's B1 security level. An example CMW is the HDS ViewStation terminal [HDS].

The SIAR-IA architecture [SIARA00] accomplishes separation without using a secure kernel. When an application requires access to trusted services, requests are made via a secured boundary interface. The application must be authenticated by the boundary device, which will attempt to carry out the requested service access on behalf of the application. All trusted services are certified. When a trusted service becomes available, it registers with the services directory, and the interface, capabilities, and limitations are stated. The services layer will arbitrate and expose the trusted service capabilities to the applications.

The SafeStart approach [SafeStart] attempts to provide some secure kernel functionality by establishing a repeatedly verifiable set of trusted software on a PC. SafeStart lets the user identify some critical minimal set of software that must be present to perform certain tasks. The integrity of this set is verified each time the system boots, using a hierarchical process.

The Trusted Coprocessor approach combines a small, verifiable TCB with a compact, tamper-resistant, fully isolated hardware platform. Sensitive operations such as cryptography and access control decisions are executed only by the Trusted Coprocessor. This approach is similar to that used to implement trusted guards at enclave boundaries. However, the coprocessor approach has the advantage of tighter hardware integration with the general-purpose platform that it protects/oversees, which results in greater reliability, communications efficiency, and simplified logistics.

3.5.4 Implementation Issues

There are several challenges to supporting C-IAA processing with a TCB:

- The research community has access only to publicly available kernels, such as Linux. Proprietary operating systems, such as Windows, must be secured by the organization that owns them.
- Operating systems are constantly updated, and secure kernel development always lags behind.
- While not true of all commercial computer platforms, PC hardware is particularly difficult to secure because there is little support for security in the hardware itself.

3.5.5 Recommended Research

3.5.5.1 Operating System Guidelines

In lieu of concrete design constraints such as target platform capacity or legacy software requirements that would dictate operating systems, drivers, interfaces, etc., attention should be focused on determining guidelines and evaluation criteria for operating systems based on the a priori security requirements of C-IAA security functions. The research project would establish minimum requirements for C-IAA operating systems and methods for evaluating and selecting candidate systems and propose methods for testing and hardening selected candidates.

3.5.5.2 Trusted Coprocessor Experiment

Retrofitting existing application platforms with secure operating systems or patches may have an adverse impact on the cost of ownership of such platforms due to such factors as

- Code impacts to existing applications

- Performance impacts resulting from inefficient security processing
- Performance impacts due to concurrency interactions
- Higher maintenance costs for secure operating systems and patches
- Less available technical and market support for secure operating systems and patches

As a hedge against such cost impacts, the Trusted Coprocessor approach should be investigated by means of an implementation experiment and demonstration. The goal of the experiment would be to integrate an off-the-self coprocessor module and security kernel with a standard Windows platform to demonstrate the execution of representative security functions on the Trusted Coprocessor with little or no modification of the application to be monitored. Objectives of the effort would include characterization of

- The coprocessor's capacity as applied to this problem,
- The effectiveness of the monitoring and ability to constrain Window's activities,
- Unintended interactions, and
- Scaleability.

3.6 C-IAA OversightNet

3.6.1 Concept

The OversightNet is a command and control subsystem dedicated to centralizing and coordinating the IA security functions of an organizational enclave. Toward that end, OversightNet collects IA data from a variety of sources, including processing nodes and dedicated network sensors, identifies and assesses IA-significant events, correlates patterns, identifies possible attacks, analyzes threat conditions, and recommends CoAs in addition to reporting status and responding to IA administration and analysts.

Knowledge and decision-making are virtually centralized, so the OversightNet can be implemented either as a distributed collection of cooperating processes or as a monolithic process, depending on implementation efficiencies and requirements. Similarly, the operations of the organizational enclave, which the OversightNet supports, may be geographically collocated or distributed.

As such, the OversightNet represents the next generation of automated IA processing beyond the currently available network sensors and ID systems.

3.6.2 C-IA Requirements

The C-IAA OversightNet must be capable of

- Secure communications with its subordinates and component elements.
- Collection of IA-significant events and conditions.
- Correlation and detection of patterns.
- Detection of simple and complex attacks.
- Situation assessment and risk analysis.
- CoA recommendations.
- Automated response.

This study has focused on the collection of functionality that results in automated IDAR, specifically, collection, correlation, detection, assessment, and response.

3.6.3 Current State of the Art

As recounted in Section 3.3.1, all the building blocks necessary to realize the OversightNet are currently available. However, as of this writing, the various technologies and products have not been integrated into the comprehensive capability envisioned by C-IAA. Therefore, interfaces, methods, and glue logic are lacking, and some key technologies must be migrated from other fields.

Summarizing each of the major areas:

- *Data collection.* A variety of standalone host- and network-based sensors are available. Few systems have attempted integration of the sort seen in network management systems.
- *Correlation.* Available products emphasize simple pattern and trend recognition. Currently, the leading work in continuous correlation is in the field of fraud detection for telephone and credit card abuse.
- *Situation and threat assessment.* Much of this work is still in research and development. Leading current examples have been built to support weapons systems.
- *Automated response.* Although technologies for self-adaptation have been used in experimental systems and deployed for the control of certain communications systems, the IA community has been reluctant to fully implement automated response capabilities.

3.6.4 Implementation Issues

Collection. Sensor subsystems continue to suffer from false positives and false negatives, thereby increasing the volume of event data while reducing its net worth due to inaccuracy. IA analysts historically either set sensors to a low threshold because they don't trust the sensor's ability to filter benign events or to a high threshold to minimize the volume of data. Neither approach leads to good attack detection, and both should be avoided in a C-IAA implementation.

Correlation. Current products have incorporated an insufficient amount of expert knowledge methods to allow the OversightNet to substantially offload the work of the IA analyst. Because the threat changes rapidly, both in response to technology growth and in direct response to better IA countermeasures, C-IAA must adapt to be successful in the medium to long run. Correlation methods must be extended to incorporate automated pattern learning and prediction.

Situation and threat assessment. The least mature of OversightNet core technologies is in the area of situation and threat assessment. Much important work is still in the laboratory. Regardless, for automated assessment to be meaningful, current methods for situation and consequence modeling, both static and dynamic, must be extended.

Automated Response. The implementation of automated response faces two impediments, one technical and one institutional. The technical issue concerns the ability to make the C^2 of automated response sufficiently robust and immune from IW attacks that it can be reliably exercised with certain results. The institutional issues concern the legal and political ability to support automated response in this arena and the impact that autonomous or semiautonomous actions by the system may have on established C^2 mechanisms within the organizational enclave.

3.6.5 Recommended Research

3.6.5.1 Collection

For the OversightNet to effectively and efficiently consume IA data from a wide variety of sources, including host and network sensors, data mining and correlation tools, etc., a consistent method of data normalization methods and data formats must be established. The project should review the work of the IETF with regard to the IDXP and related protocols when considering suitable target formats. However, because the adoption of standardized formats and semantics by the existing sensors and IDSs is problematic, sufficient attention should be placed on identifying a general-form solution for the parsing and translation of existing formats into the target forms.

3.6.5.2 Correlation

A valuable research project in the area of correlation would be the experimental selection and adaptation of an existing fraud detection tool to IA. The recommended effort would survey existing fraud detection technologies and products and select one or more products for technology transfer to IA. The objective of the effort would be to gain a better understanding of fraud pattern detection methods and to determine their direct applicability to C-IAA.

3.6.5.3 Situation/Threat Assessment

To substantially advance automated situation/threat assessment, practical methods of modeling and manipulating situation data must be identified. Targeted surveys of related efforts would be used to formulate requirements for an initial situation data model, which would lead to the construction of a prototype model. In parallel, a series of attack stimulus/response scenarios would be developed, against which the prototype model would be tested and perfected. Knowledge gained from this research would directly benefit efforts to extend existing situation/threat assessment components.

3.6.5.4 Automated Response

As with other areas of C-IAA, guidelines defining permissible limitations for automated response must be developed. Similar in nature to the recommended project to determine the sensitivity of IA data under various circumstances, this research effort should undertake to develop a prototype “rules of engagement” for automated response that cover a wide variety of threat condition scenarios. The project should produce a preliminary rules-of-engagement document for evaluation and comment by cognizant parties. In addition to legal, political, intelligence, and jurisdictional issues, the effort should also consider the systemic effects of response failures as well as successes.

3.6.5.5 Policy Research

Whether collecting data, assessing the situation, or determining an automated response, the OversightNet is carrying out one or more operational policies that map states and stimuli to new states and system responses. Currently, policy creation and management falls into two broad categories. Simple policies can be expressed and maintained without the use of formal languages and can be made to adapt in the field under the control of administrators. Complex policies are expressed in formal languages and must be maintained by those trained in the art. As a result, complex policies currently lack the flexibility of simple policies as well as the simple policy's ability to benefit from direct interaction with using administrators.

While it is anticipated that the policy(s) that govern the operations of a given OversightNet implementation will be multifaceted, it remains to be determined how complex the policy(s) will be, and therefore what manner of policy composition, maintenance, and dissemination methods and tools will be necessary to support the OversightNet.

The objective of this proposed effort would be to characterize the anticipated complexity of the OversightNet policy(s) using modeling and formal methods. The resulting characterization would be used to evaluate the existing methods of policy expression to determine those policy technologies (and products) that should be pursued for C-IAA.

Because the automated functions of the OverNet and OuterNet also require the expression of policy, this research would directly benefit the forward progress of those C-IAA components.

3.6.5.6 Retained Data Management

The automated correlation and decision-making functions envisioned for the OversightNet (and for the OverNet and OuterNet as well) depend on large volumes of IA data to be readily at hand. Moreover, by its very nature, IA data is captured fairly continuously, and will, even with very efficient filtering and significance assessment, grow without bound. C-IAA must, therefore, incorporate effective mechanisms to warehouse and archive unused data without inadvertently crippling its analysis capabilities.

For systems of the scope postulated for OversightNet, many aspects of this dynamic data management remain to be investigated, including

- The rough order of magnitude of data to be maintained under active management.
- Criteria for data aging.
- Appropriate stages of warehousing prior to archiving.

- The probability that a warehoused datum will be spontaneously required.
- The effect on system response time and effectiveness due to delays that result from data retrieval from a warehouse or archive.

The proposed effort would address these issues in conjunction with a detailed review of the existing technologies/products that address bulk online data management. The appropriate candidates would then be evaluated against the resulting preliminary OversightNet requirements. This research would also directly benefit development of the OverNet and OuterNet components.

3.6.5.7 Knowledge Discovery and Data Mining

An initial literature review indicates that some research has been done in the field of knowledge extraction from ID data, but that it remains a largely unexplored topic. It remains to be determined

- What data needs to be collected for intrusion response and prevention, in particular for C-IAA decision-making.
- The appropriate formulation of data mining decision policies.
- Which off-the-shelf tools are directly applicable to C-IAA decision-making.

Data fusion needs be accomplished for

- Sensor data by the OversightNet
- OversightNet and other data by the OverNet
- OverNet and other data by OuterNet.

ID methods must be expanded to detect coordinated distributed attacks and other “low profile,” less visible, and less known attacks. Such attacks can be sophisticated and beyond the scope of regular warning channels such as a CERT advisory, because such attacks are neither obvious nor announced.

Sophisticated attacks can be detected by finding patterns in data. Currently, it is possible to detect a single, simple functional activity based on a single information source (for example, cellular phone cloning abuse based on telephone call records). Current research is moving in the direction of fusing several related data streams to determine several related functional activities (for example, examination of several types of financial records to determine money laundering). The goal is to be able to fuse information from a high number of data sources and gain insight into the functionality of attackers with a low false alarm rate and high true positive alarm rate. Incorporation of the new generation of KDD tools would allow for automated, data-driven extraction from primary sources, automated trigger generation, automated link discovery from secondary sources and extracted facts, structural and temporal pattern learning for new situations, and performance improvement.

3.7 C-IAA OverNet

3.7.1 Concept

It is tempting to think of the OverNet as a “super OversightNet” that does for organizational enclaves exactly what the OversightNet does for individual processing platforms and their networks. While the OverNet is expected to use many of the same technologies and functions identified for the OversightNet, there are many subtle distinctions that arise from the OverNet’s position outside of the organizational enclaves, and its potential to serve several otherwise unrelated organizations. As the result of these distinctions, the OverNet is similar, but far from identical, to the OversightNet concept.

The OverNet’s role in C-IAA is to

- Coordinate and centralize situation and threat assessment across multiple organizational enclaves.
- Develop and promulgate to its participants a common IW battle view.
- Coordinate enclave responses to threat conditions and attacks.

Like the OversightNet, the OverNet may be implemented either monolithically or distributed, and the OversightNets that it supports are expected to be geographically remote from the OverNet and from each other.

3.7.2 C-IA Requirements

Like the OversightNet, the Overnet must be capable of

- Secure communications with its subordinates and component elements.
- Collection of IA-significant events and conditions.
- Correlation and detection of patterns.
- Detection of simple and complex attacks.
- Situation assessment and risk analysis.
- CoA recommendations.
- Automated response.

Additional functionality required of the OverNet includes

- The ability to provide the OversightNet services on a collaborative basis.
- Tolerance of incomplete or obfuscated data presented by subordinates.
- Tolerance of full or partial noncompliance with IA C² on the part of subordinates.
- Development and dissemination of a common battle view for its domain.

Because the OversightNet functions were addressed in association with that component, this study has focused on the additional OverNet functionality.

3.7.3 Current State of the Art

While no exact analog of OverNet functionality exists today, a few large-scale special-purpose systems have been developed, and fielded, with similar features. Examples of existing command and control and intelligence systems include

- Global command and control system (DISA)
- Cooperative engagement system (Navy)
- Missile warning (US Space Command)
- Submarine tracking network (Navy)

Examples of large-scale adaptive communications networks include

- NSF Partnership for Advanced Computational Infrastructure
- Automotive Exchange Network
- SNET
- Classified adaptive communications

3.7.4 Implementation Issues

Issues that are new to the OverNet are addressed in this section.

Collection. The possibility of incomplete and/or obfuscated data reporting by subordinates complicates data collection and requires the OverNet, in effect, to oversample the data so that missing data can be inferred. In addition, the potential for diversity in data formats and semantics is amplified when different organizations are supported.

Correlation. The degree to which implementation of OverNet knowledge and decision-making is either monolithic or distributed has a profound impact on the approach to correlation, as well as to its efficiency and survivability. Regardless of the potential benefits of either approach, it is not currently known if the OverNet can be implemented in a distributed (peer-to-peer) form or if a monolithic, physically centralized implementation is culturally acceptable under most circumstances. In either event, the support of multiple organizations exacerbates the issues associated with data aggregation and protection.

Situation and threat analysis. There exists the possibility that the decision-making algorithms employed at the OverNet level must accommodate additional uncertainty parameters to compensate for softer input data. Similarly, the resulting findings may be correspondingly “softer.” Conversely, the need to compensate for “soft” command and control greatly complicates situation and consequence modeling.

Automated response. Whereas the OversightNet must be prepared for a subordinate component to fail to carry out a CoA, the OverNet must be prepared for a subordinate enclave to choose not to carry out a CoA.

3.7.5 Recommended Research

The following research recommendations address the specific needs of the OverNet.

3.7.5.1 Collection

Three topics of study dominate the extension of data collection at the OversightNet level to that at the OverNet level.

OversightNet data contributions. Using existing, simpler IDAR systems as examples, it should be possible to predict the types of data contributions to be required of the participating OversightNets, as well as their frequency of submission.

Extended normalization. As IA data is consolidated, aggregated, and condensed in moving up the C-IAA hierarchy, it will be necessary to accommodate additional metadata characteristics as the data is normalized and reformatted. Example characteristics include sensitivity, classification, and uncertainty factors.

Impact of dynamic OversightNet participation. The difficulty of automatically supporting the spontaneous arrival or departure of a participating OversightNet is a matter for further investigation. As the complexity of maintaining a useful situation model is determined, it may become apparent that changes in participation are more cost-effectively handled synchronously via administrative action.

3.7.5.2 Correlation

A clear potential benefit of the OverNet is the creation of a common battle view for its domain so that the participating enclaves can gain a sense of the situation across their communities and better prepare coordinated responses. The proposed study would perform a detailed review of the existing common battle view technologies and identify candidates for possible technology transfer to C-IAA. Depending on the complexity of the available candidates and the estimated difficulty of transfer, this study could be extended to perform a prototype experiment that demonstrates the use of such technology applied to a virtual battlefield.

3.7.5.3 Situation and Threat Assessment

Two issues in this functional area are suitable for advanced research and investigation.

Policy expression, exchange, and maintenance. Multiple organizations are expected to imply disparity not only in data formats and semantics, but also in operational policies. If the intelligence community's experience in its attempt to normalize security marking procedures (which are a function of policy) is any indication, the normalization of policy expression should be addressed while the C-IAA is still being formulated.

Adaptive confidence factors. The decision-making process within the OverNet is expected to use confidence factors in addition to other measures of uncertainty. Because the OverNet is postulated to operate within an imperfect data sphere, it may be necessary or desirable for confidence to be modified over time in response to changing conditions. The impact of dynamic confidence factors on the situation model and the practice of decision-making should be investigated.

3.7.5.4 Automated Response

Just as the data available to the OverNet may be imperfect, so too is its command and control of the participating OversightNets. In this regard, the potential benefit of applying uncertainty techniques to C^2 theory should be investigated using formal mathematics and modeling techniques

3.8 C-IAA OuterNet

3.8.1 Concept

C-IAA conceives the OuterNet as containing one or all of the following three IA service suites:

- An OverNet of OverNets that provides the same set of OverNet functions to a collection of participating OverNets as an OverNet provides to its participating OversightNets.
- A uniform interface for law enforcement agencies to support computer emergency response events; this should support the secure interchange of IA data, forensics data, and other electronic artifacts.
- A platform that supports the widespread monitoring of public communications infrastructures to detect large-scale and subtle attacks.

Unlike other C-IAA abstract components, the OuterNet should be implemented using a separate physical network to directly reduce vulnerabilities.

3.8.2 C-IA Requirements

OuterNet requirements that are in addition to OverNet requirements previously considered include

- Collection, correlation, and assessment of public IA events.
- Secure extraction of IA data from organizations and OverNets.
- Extensive computer forensics analysis capabilities.
- Ability to recommend protective countermeasures.

3.8.3 Current State of the Art

There are no existing networks or systems resembling OuterNet. The US federal government has contemplated a government-only intranet, called Govnet, that would use a physical transport separate from the Internet. [Govnet, CDT01]. The administration issued a Request for Information on a proposal to segment off communications of confidential information among government agencies. The reaction to this proposal was not favorable: “The administration has said that it will ensure that public information would not be stored permanently on this new system, dubbed GovNet. CDT is concerned that the resources needed to develop a GovNet, combined with the security risks that would need to be addressed, would be too heavy a drain on the government’s already taxed information security and openness projects to justify the potential benefit” [CDT01].

Otherwise, the state of the art with regard to emergency response support is characterized as follows:

- Collection is typically post-event and only semiautomated.
- Correlation is achieved by analysts using non-real-time forensics tools.
- Situation/threat assessment is performed manually using forensics reports and experience.
- Automated response is scrupulously avoided.

3.8.4 Implementation Issues

3.8.4.1 Communications

To assess the feasibility of implementing the OuterNet using a physically separate transport network implementation, this study has considered a variety of alternative configurations including IP-based VPNs, circuit-based VPNs, and dedicated private transmission facilities, as well as representative costs for each. The results of this feasibility analysis are presented in the following subsections.

3.8.4.1.1 Internet Structure

The current Internet consists of the national backbone networks connected by exchange points, with numerous regional and local networks connected to the backbone. Some regional networks, such as the Metropolitan Area Ethernets (MAE) networks can be quite large.

The physical backbone network is a high-volume physical network provided by long-distance exchange carriers. A national backbone network segment is formed when a national Internet backbone provider company leases the physical network and uses it to connect high-speed routers in various locations. There are only between 20 and 30 national backbone providers, such as CompuServe, MCI, DIGEX, IBM, Sprint, ANS, and BBN. All the nodes owned by a national backbone provider are called Points of Presence (PoPs).

The key exchange points on the Internet are called Network Access Points (NAPs), with major exchange points shown in Table 5. There are four official NAPs and four de facto NAPs. There are two Federal Internet Exchange Points (FIX), used to connect MILNET, NASA Science Net, and other federal government networks. It is planned to reroute traffic from FIXs to NAPs. The Commercial Internet Exchange (CIX) points were formed in the early days of the Internet to provide routing for commercial purposes. CIX points are not used extensively today.

Table 5 — Main Exchange Points on the Internet

Function	Name	Location	Operator/owner
Official NAP	NAP	San Francisco	Pacific Bell
Official NAP	NAP	Chicago	Bellcore and Ameritech
Official NAP	NAP	Pennsauken, NJ	Sprint
Official NAP	MAE-EAST, MAE-EAST+	Washington, DC	Metropolitan Fiber Systems
Defacto NAP	MAE-WEST	San Jose, CA	"
Defacto NAP	MAE-LA	Los Angeles, CA	"
Defacto NAP	MAE-DALLAS	Dallas, TX	"
Defacto NAP	MAE-CHICAGO	Chicago, IL	"
Historical legacy NAP	FIX-EAST	College Park, MD	University of Maryland
Historical legacy NAP	FIX-WEST	Moffet Field	NASA Ames Research Center
Historical legacy NAP	CIX	Santa Clara, CA	Willtel
Historical legacy NAP	CIX	Herndon, VA	"

National backbone providers connect with each other through NAPs, PoPs, and any other exchange point feasible. Regional and local Internet service providers connect their networks to the backbone via POPs. Existing telephone lines are also used often.

3.8.4.1.2 Internet Routing

Routing of messages on the Internet is based on packet switching. In essence, packet switching can be described as dividing a message into pieces and sending each piece via a possibly different route. This approach is different from the routing used in telephone communications, where a dedicated circuit, or a consistent/persistent route, is established for the purpose of communication between two users.

Communication between two hosts on the Internet is handled by protocols at several layers. On each computer, each protocol layer communicates with the protocols of the same layer on other computers. This exchange is accomplished by passing messages with the protocols at one layer below and the layer above on the local computer, as shown in Figure 4.

When a user attempts to communicate with a user on another host, the user must invoke the appropriate higher-level protocol that will construct a message. In the OSI architecture, message construction happens at the application layer. Similarly, in the TCP/IP model, message construction happens at layers above the TCP layer. Each message is possibly divided into packets that can be carried by the network. Each packet contains header information and data. Header information conveys information needed for routing and the protocol used, such as source and destination IP, or the total number of

pieces that comprise a message. Each protocol layer adds its own header as it passes a message to the next layer down, as shown in Figure 5. When a layer receives a message, it strips the corresponding header for its layer peer.




Figure 4 — The ISO Layered Protocol Model

The message is passed to the network layer, which is responsible for routing. At the network layer, the unit of data exchanged is called a packet. Therefore, all routers and all network nodes must have at least up to network-layer functionality. The packet is passed to the data link layer, transformed into a stream of bits, and sent on the channel. The physical layer handles the transmission of bits over a communication link. In the TCP/IP model, IP is responsible for network functionality, while TCP is responsible for reorganizing packets into a coherent data stream.

On the receiver side, the same process is repeated in reverse. The physical layer detects waveforms and delivers a stream of bits. The data link layer is responsible for collecting a stream of bits into a frame and correcting transmission errors. In practice, this functionality is implemented using a network adaptor that is responsible for framing of bits, error detection, and media access protocol. The packet is passed on to the next higher layers.




Figure 5 — Protocol Encapsulation using TCP/IP and Telnet as Examples

A router is a device that works at the network layer and routes packets dynamically. Each router has dynamically updated routing tables and determines the next router to which the packet should be sent.

There are several misleading facts based on current marketing practices. Switches and routers are frequently mixed up for marketing purposes. For an informal explanation of the current confusion about routers and switches, see [Coh97]. Many Internet Service Providers (ISPs) advertise “dedicated link to the Internet,” as opposed to a dial-up connection. “Dedicated link” means that the connection to the Internet is guaranteed 24 hours per day, and a permanent Internet address is assigned to the host. However, once a packet is on the Internet, packet switching is used for routing and there is no guarantee of performance.

3.8.4.1.3 Internet Connection Costs

The Internet at the physical level requires point-to-point circuits between routers. The cost of laying fiber-optic or copper cable includes the material, construction, and right-of-way. Fiber-optic cable is the most expensive but most efficient networking medium and has a lower installation cost than copper because of its small size and lightweight. Fiber is also more secure, because it is currently difficult to tap. In this report, we will discuss creating a physically separate OuterNet based on fiber-optic cable.

Laying fiber-optic cable in an urban environment is much costlier than in rural areas. Some practical estimates are \$16,000 per mile for rural areas and \$500,000 per mile for a city business district. However, some contractors claim that it could be done for under \$10,000 per mile in rural areas [ISP]. FCC’s Benchmark Cost Proxy Model (BCPM) contains an official guide to the cost of laying fiber. It is quite popular to install cable along highways and railroad beds, which also brings the installation cost down [Lev3, ZD].

The right-of-way costs can be very high. States can charge high taxation cost for installing cable along interstate highways. Railroad companies also demand fees. For example, sample taxation proposals from Utah, Arkansas, Florida, Kentucky, and Minnesota include a one-time \$500-per-mile charge, an annual fee of \$1,000 per mile, and even considerations for tax of \$250,000 per mile. The most current laws, legislations, and proposals need to be taken into account when considering laying cable.

The following costs, except construction fees, would be approximately similar in the case of a separate physical connection or a virtual network. Possibly high construction fees, as outlined above, would be involved in laying down a separate physical connection. According to [Dowd96], some representative connection costs are given in Table 6.

Table 6 — Sample Internet Connection Costs

Item	Payment schedule	Notes
Network termination unit	One-time	
Channel Service Unit (CSU)/Digital Service Unit (DSU) or Multiplexer	One-time	
Special construction fees	One-time, if applicable	Can be high, if it involves laying cable
Circuit installation	One-time	
Circuit fees	Monthly, for leased lines	
Public data network install	One-time, if applicable	
Flat public network fee	Monthly, if applicable	
CIR fee	Monthly, if applicable	
Port install	One-time, if applicable	
Port fee	Monthly, if applicable	
Internet service install	One-time	
Internet service	Monthly	

Sample equipment and connection pricing, quoted from [Dowd96] p.6, are given in Table 7.

Table 7 — Representative Connection Costs

Item	Cost (\$)			
	56Kbit/sec frame relay	T1 dedicated	10Mbit/sec SMDS	T3 via NAP
CSU/DSU or Multiplexer	250	1,000	3,000	5,000
Router	1,800	1,800	2,500	6,000
Server	2,500	2,500	3,500	10,000
Transport one-time	800	1,200	3,000	5,000
Transport recurring annually	1,500	4,800	36,000	25,000
Internet service one-time	1,200	2,500	3,500	5,000
Internet recurring annually	4,800	18,000	30,000	60,000
Total, year one:	12,850	31,800	81,500	116,000

According to [Net2K2NE], the approximate cost is shown in Table 8.

Table 8 — Representative Connection Costs

Item	Cost (\$)	
	56Kbit connection	T1 connection
CSU/DSU	350.00 * 2	900.00 * 2
Router	800.00	900.00 to \$2,000.00
Startup configuration fee	1,000.00 [Net2K2NE] 2,950 [TeleWeb] 775 [SVS]	1,500.00 [Net2K2NE] 3,495 [TeleWeb]
Monthly fee	225 [TeleWeb] 675 [SVS]	900 [TeleWeb]
Leased line fee	Obtainable by request	

A Channel Service Unit (CSU)/Digital Service Unit (DSU) is most simply described as a leased line modem.

A CSU/DSU is needed at a customer's location as well as at the ISP's network operation center (NOC). Two units are required, one for each end of the connection. A router is required at the customer's location; a router port is required at the ISP's location to service the leased line connection. The ISP's network must be configured to accept the customer's connection and also to ensure that the customer's site is communicating compatibly with the ISP's NOC and with the Internet.

3.8.4.1.4 Leased Line versus Virtual Private Network

If a network is not constructed using a separate physical connection, it must be constructed using the existing lines, either as a dedicated leased line or a VPN. In both cases, the following costs are similar: Internet connectivity (line and adapter cost); server hardware software and licenses; installation and setup; and modem costs for clients.

In the case of a leased line connection, both CSU/DSU units remain the customer's property, and the CSU/DSU used for the ISP's side of the connection is returned to the customer at the end of the contract term. The customer is responsible for maintaining the connection and paying the telephone charges to connect to the leased line. The telephone charges for remote access to the leased line, as well as the modems for the remote users, can be very high. In the case of a VPN, the ISP manages the connection, and the phone line charges are eliminated. A comparison in cost is shown in Table 9 and Table 10 [SVS].

The data in Table 9 and Table 10 is based on the assumption that, on average, each user will spend 12 hours per month connected to the network, the average number of connected users at any given time will be 15, and the cost for long-distance charges averages \$.13 per minute.

Table 9 — One-time Setup Costs

Item	Cost (\$)
Phone line group setup (20 lines @ \$20/each)	400
Modem bank (20 ports @ \$100/each)	2,000
Total (one-time cost)	2,400

Table 10 — Monthly Recurring Costs

Item	Leased line cost (\$)	VPN Cost (\$)
Phone line maintenance (20 lines @ \$15/month/line)	300	N/A
ISP charges for Internet access (50 total users)	N/A	1,000 (20 per user, unlimited use)
Long distance charges ¹	1,404	N/A (assumed local dialing)
Total (per month)	1,704	1,000

Therefore, leased lines provide a cost-effective solution if supporting few users with minimal long-distance and data transport charges. However, a VPN can be a more cost-effective method of choice for supporting more users. Some hidden costs of VPN use include the fee paid to the ISP to maintain client access, as well as compatibility issues in relatively new and untested VPN technology.

3.8.4.2 Collection

Issues associated with the collection of public access networks for the purpose of monitoring those networks are subject to both legal and political considerations. Assuming it were legal to collect such data at, for example, Internet Exchange Points, secondary legal issues would arise with regard to the ownership and safeguarding of intellectual property and similar matters. Political issues are exemplified by the concern of sensors being overt versus covert and the need to establish a new classification channel for such surveillance if covert.

Due to the potential for overwhelming quantities of IA data collected, new extended strategies for selective collection durations and places would have to be developed.

3.8.4.3 Automated Response

Automated response in the realm of the public access networks is anticipated to be infeasible due to additional sensitive legal and political issues.

Otherwise, automated response directives cascading down to participating OverNets to their participating OversightNets can use many of the same interorganizational negotiations and agreements that are necessary to enable OverNet automated response.

3.8.5 Recommended Research

3.8.5.1 Communications Topology

There are several possibilities for an OuterNet topology:

- A physically separate network physically connected to the current Internet. The network would be constructed in the same way that the current Internet is constructed, by installing new cable. Considering the fact that the current Internet keeps on growing and the high cost of laying cable, it may be cost-prohibitive to attempt to physically parallel the Internet.
- A network that connects to the exchange points on the Internet.
- A separate physical network or a virtual overlay network on top of the current Internet.
- A network that connects OverNets, using either virtually or physically separate transport.

Since most of the Internet is owned by the commercial sector, it is impossible to obtain the exact topology of the current Internet. The key exchange points are known, and many PoPs are known. However, even if a PoP's exact location is not known, it is possible to know who owns it. The Internet topology is published quarterly by *Boardwatch* magazine [Board]. The number of current key exchange points is rather small, and the number of PoPs should be on the order of few hundred [Rick97]. It is necessary to investigate which exchange points need to serve as monitoring points.

Example analysis using a hypothetical approach. The OuterNet is a separate network that has nodes at NAPs. An OuterNet node consists of a packet sniffer, such as Snort, and an OuterNet module. An OuterNet module collects statistical information necessary for its OuterNet decision-making. It would be desirable that the OuterNet has modules at PoPs as well, if possible. Since the locations of some PoPs are not known, but it is known which organizations own all PoPs, the OuterNet could establish a presence at all PoPs. Most packets traveling across networks will go through NAPs and PoPs. However, some packets will not need to pass through exchange points, and it would be impossible to monitor such traffic unless a physical copy of the Internet is constructed

from each POP down. We estimate that most traffic will not fall into the latter category, and thus the cost–benefit ratio justifies the OuterNet nodes at NAPs and PoPs only.

To capture traffic that possibly did not pass through an exchange point, each OuterNet module also receives input from OverNet modules that are closest to it. Also, each large Web site host such as yahoo.com can decide to install an OuterNet module.

The OuterNet network duplicates the Oversight/OverNet relationship. The OuterNet module is based on the same principles as OverNet and OversightNet modules, the only difference is the data used for decision-making and the policy used. The “Oversight” portion of the OuterNet consists of decision modules that reside at exchange points and classify packets passing through in real time as well as over the long term. Groups of OuterNet modules, or possibly all OuterNet modules, are connected to a top-level OuterNet module, which acts similarly as the OverNet module to the OversightNet.

As the government operates the OuterNet, contacting the OuterNet would be similar to contacting the police and inviting them to further the investigation. Therefore, it remains to be investigated what is possible under the current law and if it is permissible to install OuterNet modules at NAPs and PoPs. NAPs were constructed under NSF grants, and PoPs are privately owned. We could argue that the OuterNet acts as a hired security guard and ask Internet subscribers to sign a statement of acceptance. We could also argue that the OuterNet network is equivalent to having traffic police on the highways. In the case of making the analogy to the traffic police, it would be necessary to determine what constitutes violations of Internet “traffic.” Activities that are illegal in the “ordinary” world are also illegal on the Internet, but in the ordinary world, citizens are not monitored by default—only if they are reported by someone or caught. It remains to be investigated what kind of information and under what circumstances can be collected by the OuterNet.

NAP and PoP owners decide to switch on the OverNet modules, and the OverNet modules decide when, if at all, to send alerts to the OuterNet module. The OuterNet module processes the information received, determines CoA, informs the OverNet module of it, and must also inform the human operator of the proposed CoA. It is assumed that the human operator must issue permission for OuterNet personnel to perform the actions proposed, such as setting up a honey-pot or monitoring traffic of the OverNet.

The legal issues related to government ownership of the OuterNet and the rights to collect data on various networks will play a significant role in OuterNet architecture and functionality. An alternative to a government-owned OuterNet is a series of private second-level OverNets owned by private companies that specialize in distributed security. These companies would be the “man in black” (MIB) that can be hired to solve security problems. An organization most analogous to this concept today is the CERT advisory at Carnegie Mellon; however, CERT simply receives reports from various sites that

volunteer to report and then disseminates the information. The OuterNet could be more proactive, because it would perform more rigorous analysis of individual sites.

3.8.5.2 Access and Rights Guidelines

It is desirable to establish the legal baseline that constrains the architecture, operations, and functionality of OuterNet implementation as soon as possible. This proposed research area would undertake the following:

- Review current enabling surveillance legislation.
- Identify current minimum and maximum legal requirements.
- Identify potential monitoring points, if permitted.
- Survey potential participants for sensitivity issues.
- Determine need for sensor mobility, if permitted.
- Determine need for covert surveillance, if permitted.

The result of this investigation would be two documents: Access and Rights Guidelines for keeping OuterNet within the current law and Implementation limitations due to current law and their impact on desired operations.

3.8.5.3 Other Research

Topics supporting OuterNet development that should be addressed as time permits:

- Investigate methods for the normalization of forensics data and select high-potential candidates for use in C-IAA.
- Adapt OuterNet situation/threat analysis to consume forensics results as IA input data.
- Estimate the volume of events to be handled by the OuterNet so that performance and scalability implementation guidelines can be established.

3.9 Decision-making in C-IAA

This section presents a preliminary design concept for automated security response on all levels from processing node to *OuterNet*. To aid in efficient and effective action and decision-making, as many human tasks as possible need to be automated and the pertinent information organized in a user-friendly way. This will make the process more efficient as large volumes of information pass through the network in real time. The human input is crucial for proper operation, where automation supports efficient and effective action and decision-making.

This section provides the initial approach for the following:

- Purpose of OversightNet-OverNet exchanges and the roles of each network
- Data exchanged by OversightNet and OverNet
- Protocol and authentication mechanisms for data exchange
- Actions that can be taken by the OversightNet and OverNet

The design goal is to keep implementation efficient and effective.

3.9.1 Stages and Tiers

C-IAA provides automated decision-making assistance on the OversightNet, OverNet, and OuterNet levels. Each level is called a tier.

An OversightNet performs decision-making on tier I. OversightNet receives sensor messages and classifies them according to the criticality level based on user-specified policy. Sensors can be geographically distributed. It is most likely that each sensor is monitored by only one OversightNet. Therefore, OversightNet “knows” the local state covered by the sensors that feed the OversightNet.

OverNet performs decision making on tier II. OverNet “knows” the global state defined by all OversightNets that it monitors. The OverNet network has input into the OversightNet based on the feedback loop approach. An OverNet receives the status of its monitored OversightNets, processes it according to the user-defined policy, and either requests or changes the state of OversightNets. For example, if the OverNet concludes that most OversightNets are under attack, it can request or change the security level of the OversightNets. In real-life implementations, it is more likely that the OverNet will only request the change of OversightNet state, because of implementation issues as well as policy issues regarding autonomy of individual OversightNets.

OuterNet performs decision-making on tier III. OuterNet “knows” the global state defined by all OuterNets and other networks that it monitors.

Each tier is based on the same decision-making process, but with different parameters. In each tier, decision-making is performed in two stages. The first stage pertains to real-time event classification, and the second pertains to long-term, overall patterns.

The first stage of decision-making is performed in real time and is used to sort input events into categories of maliciousness and potential threat. This stage is based on relatively sparse information contained within a single alert. It cannot see the “big picture,” the *context* of the input events.

The second stage of decision-making is performed on long-term data. It is based on pattern recognition of attack signatures. It can also take one event that is known to be malicious and use that information to determine if there are any other events related to it. For example, an SHH can trigger an alarm because the source IP is unknown. We can look through data to see if this IP was involved in other attacks and to find out if any patterns were repeated [John00].

Both stages of decision-making are performed on local OversightNets, the OverNets, and the OuterNet. On the OversightNet, events are sensor messages (called “alerts”); therefore, the long-term overall patterns are based on data mining of sensor alerts. On the OverNet, events are outputs of the local OversightNets; therefore, long-term overall patterns are based on outputs of all OversightNets that this OverNet monitors. OuterNet input events are outputs of all monitored OverNets, as well as other inputs. Decision-making on the OuterNet will be discussed separately.

Regardless of the exact form of the input to the decision-making process, the process is in essence the same. The first-stage decision-making process is based on fuzzy classification, and the second-stage decision-making process is based on fuzzy inference rules. Fuzzy classification is used to sort incoming events in real time into predefined categories. Fuzzy inference rules are used to extract complex, more qualitative conclusions based on heuristics and quantitative information. The main difference between OversightNet and OverNet decision-making is in the data collected, information extracted, and policy used for decision-making.

Information collected must support detection of malicious events, management oversight, and response to an event. The CoA suggested must take into account the following: which responses the organization is capable of performing, what data is needed for each response, and how the data can be collected. It is not advisable to overtly collect information, for example, by using DNS or PING, because such action would allow the attacker to determine what measures are in place and what the thresholds are [John00].

Decision-making on the OversightNet and the draft design of the OverNet decision-making are presented in the following sections.

3.9.2 OversightNet Decision-making Example

In this section, decision-making process will be illustrated using the example of the OversightNet. The process is the same for OverNet and OuterNet, but the input and output events and the policy are different, as specified in Section 3.9.1.

Decision Engine (DE) modules perform decision-making in CIAA. DE is a forward-chaining rule-based expert system. The rules are defined in a policy. A policy consists of a set of IF-THEN rules. Each rule consists of two parts: the antecedent (or the conditional part, consisting of logical predicates, on the left-hand side) and the consequent (or right-hand side). If an antecedent is found to be satisfied by actual data, the consequent of the rule is asserted to hold, and the rule is said to be activated and to have fired.

The DE stage consists of two sub-modules: DE Stage I and DE Stage II.

The DE Stage I module performs Stage I decision-making. On OversightNet, it accepts input from multiple sensors and classifies each sensor alert in real time. The classification relies on various system resources, such as a list of hotlisted IPs. The information taken into account for DE Stage I decision-making on OversightNet has been discussed earlier. DE Stage I is designed for quick initial classification of each incoming alert in real time, so response to critical alerts can be immediate.

However, many long-term attacks might be inconspicuous and will not raise critical alerts. Therefore, decision-making on the next level of abstraction requires examination of statistical data patterns over a sufficiently long period of time, as well as an overview of the current situation and risks, both local and remote. DE Stage II decision-making modules provide this big picture view. DE Stage II consists of the Situation Assessment (SA) and Risk Analysis (RA) modules. SA and RA use long-term data to report on patterns of misuse and estimate the current situation, actions to be taken, and potential risks, as outlined in Section 3.3.3. The Threat Monitoring (TM) module performs data mining and statistical analysis on long-term data and attempts to detect and describe attacks in progress. The Local OversightNet Control (LONC) module consists of a TM module and the data cache used to store sensor alerts.

3.9.3 Decision Engine Stage I

The DE Stage I module helps security analysts find the most critical sensor alerts that need immediate attention and discard the false alarms. Sensor alerts are classified in real time into predetermined categories, using a user-specified policy. Sample categories are Critical, Serious, Routine, and Clear. A CoA is also suggested. A sample policy might be

```
IF source IP is hotlisted OR message is sent at night
  THEN alert is Critical, CoA is to notify supervisor.
```

3.9.4 Decision Engine Stage II

DE Stage II contains SA and RA modules and makes decisions on long-term data for broad geographic/organizational distribution, using more qualitative parameters. DE Stage II relies on TM modules for data mining and statistical data processing. Recall that DE Stage I receives alarms from sensors and individually classifies them in real time. The alarms are stored in a data cache, and the TM module performs additional statistical processing. DE Stage II receives input from the TM module and performs decision-making using SA and RA policies. SA and RA have separate policies. Both policies have the same general format and are processed in the same way, but the rules are different because the modules make decisions using different data.

The inputs to the SA module, i.e., the data that the SA policy takes into account, are system variables and TM output. SA module displays the following output to the user:

- Situation assessment and confidence in it
- Recommended CoAs

In addition to the SA output obtained from the policy table, SA displays a text description of the attack, provided by the TM.

RA policy takes into account SA's output, all values used in making SA's decision, and additional variables. RA selects the best CoA and estimates risk associated with executing this CoA.

The general format of SA/RA policies, data categories to be taken into account, and the possible values of data are shown below:

- Data Category 1: Fuzzy, Crisp
- Data Category 2: Fuzzy, Crisp
- Data Category N : Fuzzy, Crisp
- Threat Presence: Integer, 0,...,10
- Threat Level: Integer, 0,...,10
- Situation Assessment: Emergency, Critical, Suspicious, Routine
- Confidence Factor (SA): Real Number, 0,...,1
- Recommended CoA: $CoA_1, \dots, CoA_N$
- Data Category $N+1$: Fuzzy, Crisp
- The Best CoA: One or More of Recommended CoAs
- Risk Assessment: Integer, 0,...,10
- Confidence Factor (RA): Real Number, 0,...,1.

Assume that data categories 1- N are used for SA decision-making, and data categories beyond N are used for RA decision-making. Data categories have a user-defined range of values, which can be either crisp or fuzzy. Crisp values are the "usual" numerical values.

Fuzzy values are overlapping descriptive labels that are processed according to fuzzy logic principles. Sample fuzzy values might be Low, Medium, and High, with 25% overlap.

Recall that DE is a forward-chaining, rule-based expert system. DE Stage II incorporates fuzzy inference rules, as specified above. In addition, each rule is assigned a confidence factor (CF), asserting the confidence that the rule is correct. Each incoming piece of data is assigned a CF, asserting that the data value is correct. Attaching a CF to each data field is necessary because data is often estimated or derived from another set of rules. For example, Threat Level is estimated using a separate policy, described in the next section.

Data Categories 1 through N will be defined in future work. A sample SA policy is presented below, with SA factors in Table 11 and RA factors in Table 12.

Table 11 — Sample Policy with SA Factors

Rule	Current THREAT -CON	Current Political Situation	Threat Presence	Threat Level	Situation Assessment	CF(SA)	Recommended CoAs
1	High	Bad	High	High	Emergency	1	Call men in black Shut down access to network Shut down host
2	NOT High	NOT Bad	High	High	Critical	0.9	Call staff member X, Shut down host
3	Low	Good	High	Med	Serious	0.8	Email staff member Y

Table 12 — Sample SA Policy with RA Factors

Rule	Host Usage	Best CoA	Risk Analysis Conclusion
1	-Server -Mission critical -Low -Other	All Call All All	High High Low Low
2	-Server -Mission critical -Low		
3	-Server -Mission critical -Low	All	

The OverNet module performs decision-making on the OverNet. The OverNet module performs decision-making in the same manner as a DE Stage II module, as outlined in Section 3.9.4, using the same format of policy tables. However, the meaning of data fields is different, and Threat Level is calculated differently.

4 References

- [ActiveNets] Agent Based Systems Research,
<http://www.darpa.mil/ito/ResearchAreas/ActiveNetsList.html>.
- [AdTu01] G. Adomavicius, A. Tuzhilin, "Expert-Driven Validation of Rule-Based User Models in Personalization Applications," *Data Mining and Knowledge Discovery* 5, Kluwer Academic Publishers, The Netherlands, 2001, pp. 33-58.
- [ALA01] R. Au, M. Looi, P. Ashley, "Automated Cross-Organisational Trust Establishment on Extranets," *Proceedings of the Workshop on Information Technology for Virtual Enterprises*, ITVE 2001, pp. 3-11.
- [Alen00] J. Alen et al., "State of the Practice of Intrusion Detection Technologies," *Networked Systems Survivability Program*, Carnegie Mellon Software Engineering Institute, Report CMU/SEI-99-TR-028, ESC-TR-99-028, January 2000,
<http://www.sei.cmu.edu/publications/documents/99.reports/99tr028/99tr028title.html>.
- [AlFo98] J. Alves-Foss, "Multiprotocol Attacks and the Public Key Infrastructure," *Proceedings of the 21st National Information Systems Security Conference*, Arlington, VA, October 1998, pp. 566-576.
- [ANX] Automotive Industry Action Group home page, <http://www.aiag.org>.
- [ArcotID] Arcot Systems home page, "Secure Digital Identities: theArcotID,"
<http://www.arcot.com/pdf/The%20ArcotID%20v2.pdf>.
- [ArcotVPN] Arcot Systems home page, "ArcotTM for VPN," <http://www.arcot.com/vpn.html>.
- [Atki97] R.J. Atkinson, "Towards a More Secure Internet," *Computer*, vol.30, Jan. 1997, pp. 57-61.
- [BEEP] R. Cover, *Blocks eXtensible eXchange Protocol Framework (BEEP)*,
<http://xml.coverpages.org/beep.html> (modified September 2000).
- [BEEP01] *The Block Extensible Exchange Protocol (BEEP) Core*, RFC3080, July 2001,
<http://www.cis.ohio-state.edu/cgi-bin/rfc/rfc3080.html>.
- [Beha99] M. Berthold, D.J. Hand, *Intelligent Data Analysis: an Introduction*, Springer, 1999.
- [Bitta01] M.D. Bitta, "Airtight Linux," *PC Magazine*, June 12, 2001.
- [Blei98] D. Bleichenbacher, "Chosen Cyphertext Attacks Against Protocols Based on the RSA Encryption Standard PKCS#1," *Proceedings of CRYPTO '98*, August 1998, pp. 1-12.

- [Boar98] *Boardwatch* magazine press release, "Third keynote/Boardwatch Index of Backbone Providers," 1998, <http://www.keynote.com/news/announcements/pr031898.html>.
- [Boar01] *Boardwatch* magazine's "Directory of Internet Service Providers," 13th edition, Penton Publishing, Spring 2001, <http://www.ispworld.com>.
- [Bora01] S. Boran, "Top down approach to improving security," *IT Security Cookbook*, under Open content License, 2001, <http://www.boran.com/security/IT1x-2.html#Heading12>.
- [BuTh99] K.E. Burn-Thornton, S.I. Thorpe, "Improving Clinical Decision Support using Data Mining Techniques," *Proceedings of the SPIE Conference on Data Mining and Knowledge Discovery: Theory, Tools, and Technology*, 1999, SPIE vol. 3695, 0277-786X/99, pp. 207-214.
- [BVG98] K. Bansal, S. Vadhavkar, A. Gupta, "Brief Application Description: Neural Networks Based Forecasting Techniques for Inventory Control Applications," *Data Mining and Knowledge Discovery 2*, Kluwer Academic Publishers, The Netherlands, 1998, pp. 97-102.
- [BW00] R. Butler et al., "A National-Scale Authentication Infrastructure," *Computer*, 2000, pp. 60-66.
- [CaseR] i2 home page, *CaseRunner* tool, <http://www.caserunner.com/flash/product/index.htm>.
- [CC] Common Criteria Implementation Board (CCIB), *The Common Criteria for Information Technology Security Evaluation*, vs.2.1., January 2000, <http://csrc.nist.gov/cc>.
- [CCL00] L.L.H. Chung, K.C.C. Chan, H. Leung, "Discovering Fuzzy Clusters in Databases Using an Evolutionary Approach," *Proceedings of SPIE vol. 4057 (2000) In Data Mining and Knowledge Discovery: Theory, Tools, and Technology 2*, 2000, 0277-786X/00, pp. 11-21.
- [CDT01] Center for Democracy and Technology, Access to Government Information, "CDT Comments on US GovNet Proposal," November 29, 2001, <http://www.cdt.org/righttoknow>.
- [CIDF] Common Intrusion Detection Framework (CIDIF).
- [CKSS00] G. Caronni, S. Kumar, C. Schuba, G. Scott, "Virtual Enterprise networks: The next generation of Secure Enterprise networking," *Proceedings of the 16th Annual Conference on Computer Security Applications, ACSAC'00*, 2000, pp. 42-51.
- [CoAA95] F. J. Cooper et al., *Implementing Internet Security*, New Riders Publishing, MI, 1995.
- [Cohe97] J. Cohen, "Do you wanna buy a switch?" *Network World*, August 25, 1997, <http://www.nwfusion.com/news/1997/0825switch.html>.

- [CoPr01] C. Cortes, D. Pregibon, "Signature Based Methods for Data Streams," *Data Mining and Knowledge Discovery*, no. 5, 2001, pp. 167-182.
- [CSAP21] T.R. Metcalf, "Computer Security Assistance Program for the Twenty-First Century (CSAP21), Functional Requirements – Draft," MITRE technical Report, controlled distribution, Contract F19628-94-C-0001, project 039804536C, Bedford, MA, December 1997.
- [DARPA] Defense Advanced Research Agency (DARPA) home page, <http://www.darpa.mil>.
- [DCP00] V. Dhar, D. Chou, F. Provost, "Discovering Interesting Patterns for Investment Decision-making with GLOWER – A Genetic Learner Overlaid with Entropy Reduction," *Data Mining and Knowledge Discovery 4*, Kluwer Academic Publishers, The Netherlands, 2000, pp. 251-280.
- [DeFe98] L. Delgrossi, D. Ferrari, "The Design of Supranet Security Mechanisms," *Proceedings of the 7th IEEE Intelligent Network Workshop, IN'98*, 1998, pp. 167-173.
- [DeWitt] D. DeWitt, "Machines in the Myths: The state of Artificial Intelligence," ChipCenter Web page, <http://www.chipcenter.com/columns/ddewitt/col002.html>.
- [DGG99] N.G. Duffield, P. Goyal, A. Greenberg, "A Flexible Model for Resource Management in Virtual Private Networks," *SIGCOMM'99*, Cambridge, MA, 1999, pp. 95-108.
- [DiDi00] J. Dickerson, J. Diskerson, "Fuzzy network Profiling for Intrusion Detection," *Proceedings of International Conference on the North American Fuzzy Processing*, July 2000, pp. 301-306.
- [Dodg00] M. Dodge, "Atlas of Cyberspace," Cyber-Geography Research, Center for Advanced Spatial Analysis, University College, London, http://www.geog.ucl.ac.uk/casa/martin/atlas/more_isp_maps.html.
- [Dowd96] K. Dowd, *Getting Connected: The Internet at 56K and Up*, O'Reilly and Associates, 1996.
- [DSM98] L.R. Dondeti, A. Samal, S. Mikherjee, "A Dual Encryption Protocol for Scalable Secure Multicasting," *Proceedings of the Fourth IEEE Symposium on Computers and Communications*, 1998.
- [EELD] T. Senator, "Evidence Extraction and Link Discovery (EELD)," DARPA ISO briefing for BAA01-27, June 2001, http://www.darpa.mil/iso2/EELD/EELD_BAA.ppt.
- [Entrust] Entrust, Inc. home page, "Digital Signatures - Best Practice for e-Business Transactions," <http://archive.entrust.com/digitalsig/whitepapers.htm>.

- [ESBP96] M. Esmaili, R. Safavi-Naini, B. Balachandran, J. Pieprzyk, "Case-Based Reasoning for Intrusion Detection," 1063-9527/96, IEEE Press, 1996, pp. 214-223.
- [FaPr97] T. Fawcett, F. Provost, "Adaptive Fraud Detection," *Data Mining and Knowledge Discovery I*, Kluwer Academic Publishers, The Netherlands, 1997, pp. 291-316.
- [Fern01] F. Fernandez, Director of Defense Advanced Research Projects Agency (DARPA), before the Subcommittee on Emerging Threats and Capabilities, Committee on Armed Services, United States Senate, March 2001.
- [FIAC01] The First Annual Federal Information Assurance Conference (FIAC), the University of Maryland, College Park, MD, 2001, <http://www.fbcinc.com/FIAC>.
- [FlAt00] C. Flack, M.J. Atallah, "Better Logging through Formality: Applying Formal Specification Techniques to Improve Audit Logs and Log Consumers," Springer-Verlag Berlin Heidelberg 2000, RAID 2000, LNCS 1907, pp. 1-16.
- [FTC-29] R.A. Maxion, organizer, M. Dacier, S. Saydjari, guests, "Special Seminar on Intrusion Detection," *29th International Symposium on Fault Tolerant Computing (FTC-29)*, June 1999.
- [FuSh00] L.M. Fu, E. H. Shortliffe, "Application of Certainty Factors to Neural Computing," *IEEE Transactions on Neural Networks*, vol. 11, no. 3, May 2000.
- [Gass88] M. Gasser, *Building a Secure Computer System*, Van Nostrand Reinhold Company, Inc., NY, 1988.
- [GoJe98] M. Goldszmidt, D. Jensen, editors, "DARPA Workshop on Knowledge Discovery, Data Mining and Machine Learning (KDD-ML): Recommendations Report," DARPA Workshop on KDD-ML, Carnegie Mellon University, June 1998, [http://www.darpa.mil/iso2/EELD/KnowledgeDiscovery_DM_andML_Report_for_EELD_\(2\).doc](http://www.darpa.mil/iso2/EELD/KnowledgeDiscovery_DM_andML_Report_for_EELD_(2).doc).
- [GovNet] E. Lundquist, "A Call to Arms to Build Supersecure Govnet," *eWeek* magazine home page, posted October 15, 2001.
- [Grot00] R. Groth, "Data Mining: Building Competitive Advantage," Prentice-Hall, 2000.
- [GuMo98] V. Gupta, G. Montenegro, "Secure and Mobile Networking," *Mobile Networks and Applications*, vol. 3, 1998, pp. 381-390.
- [HaMa] C. Hagn, W.H. Markwitz, "Mobile Teleworking: Some Solutions and Information Security Aspects," *Proceedings of the EUROCOMM 2000, Information Systems for Enhanced Public Safety and Security*, IEEE/AFCEA, 2000, pp. 322-325.

- [HDS] HDS ViewStation System administrator's Guide,
http://www.math.psu.edu/local_doc/hds/html/3-12-6.html.
- [Hers99] E. Herscovitz, "Secure Virtual Private Networks: The Future of Data Communications," *International Journal of Network Management*, vol. 9, 1999, pp. 213-220.
- [HKMY98] Y. Hamuro, N. Katoh, Y. Matsuda, K. Yada, "Mining Pharmacy Data Helps to Make Profits," *Data Mining and Knowledge Discovery 2*, Kluwer Academic Publishers, The Netherlands, 1998, pp. 391-398.
- [HNC] HNC home page, <http://www.fraudconsulting.com/>.
- [HOSANA] Arizona State University, "Highly Structured Architecture for Network Security (HOSANA) Project," <http://www.cs.arizona.edu/xkernel/hpcc-blue/linux.html>.
- [HURD] Free Software Foundation, GNU HURD: Free Replacement for UNIX kernel,
<http://www.gnu.ai.mit.edu/hurd>.
- [IDMEF] D. Curry, H. Debar, "Intrusion Detection Message Exchange Format Data Model and Extensible Markup Language (XML) Document Type Definition," IETF Intrusion Detection Working Group, draft-ietf-idwg-idmef-xml-04.txt, September 18, 2001 (expires: March 17, 2002), <http://www.ietf.org/internet-drafts/draft-ietf-idwg-idmef-xml-04.txt>.
- [IDIP] Intrusion Detection and Isolation Protocol project (IDIP),
<http://seclab.cs.ucdavis.edu/projects/idip.html>,
<http://seclab.cs.ucdavis.edu/response/overview.html>,
<http://www.darpa.mil/ito/psum1999/J104-2.html>,
<http://www.darpa.mil/ito/psum2000/J104-2.html>.
- [IDXP] B. Feinstein, G. Matthews, J. White, "The Intrusion Detection Exchange Protocol (IDXP)," IETF Intrusion Detection Working Group, draft-ietf-idwg-beep-idxp-03, September 11, 2001 (expires: March 12, 2002), <http://www.ietf.org/internet-drafts/draft-ietf-idwg-beep-idxp-03.txt>.
- [IETFWG] Internet Engineering Task Force, "Active IETF Working groups: Security Area" http://www.ietf.org/html.charters/wg-dir.html#Security_Area (last updated September 2001).
- [IKP95] K. Ilgun, R. Kemmerer, P.A. Porras, "State Transition Analysis: A Rule-Based Intrusion Detection Approach," *IEEE Transactions on Software Engineering*, vol. 21, no. 3 0098-5589/95, 1995, pp. 181-199.
- [ImVi99] T. Imielinski, A. Virmani, "MSQL: A Query Language for Database Mining," *Data Mining and Knowledge Discovery 3*, Kluwer Academic Publishers, The Netherlands, 1999, pp. 373-408.

- [InDi] Information Discovery Web page, "A Characterization of Data Mining Technologies and Processes," <http://datamining.com/dm-tech.htm>.
- [Integr] Green Hills Software, Inc., "INTEGRITY™ Real Time Operating System," <http://www.ghs.com/products/rtos/integrity.html>.
- [ICSA] ICSA labs home page, <http://www.icsa.net>.
- [ISPPlanet] ISP-Planet home page, "The Price of Laying Fiber," http://www.isp-planet.com/business/fiber_price_bol.html.
- [Jord00] C. Jordan, "Analyzing IDS Data," May 2000, <http://www.securityfocus.com>.
- [Kasa98] N. Kasabov, *Foundations of Neural Networks, Fuzzy Systems, and Knowledge Engineering*, The MIT Press, 1998.
- [KCC00] D.J. Korsmeyer, E.T. Chow, M.P. Conroy, "IsoWAN: A NASA Science and Engineering Information and Services Network," *Proceedings of the 5th IEEE Symposium on Computers and Communications, ISCC 2000*, pp. 534-539.
- [Khal00] R. Khalnolkar, "Security is Essential to Running an E-Business," *Internet Security Advisor Magazine*, July/Aug. 2000, pp. 20-23.
- [KlDo89] P. Kline, S. Dolins, *Designing Expert Systems*, Wiley and Sons, 1989.
- [Kolu01] SSH Communications Security Corp. home page, "Kolumbus to Offer its Customers Certificates and Certificate-Based Services," press release, August 2001, <http://www.ssh.fi/about/press/detail.cfm?id=524>.
- [KoPr01] R. Kohavi, F. Provost, "Applications of Data Mining to Electronic Commerce," *Data Mining and Knowledge Discovery 5*, Kluwer Academic Publishers, The Netherlands, 2001, pp. 5-10.
- [Kosk92] B. Kosko, *Neural networks and Fuzzy Systems*, Prentice Hall, 1992.
- [Kosk93] B. Kosko, *Fuzzy Thinking*, Hyperion, 1993.
- [LaSr00] C. Labonte, S. Srinivas, "Group Management Strategies for Secure Multicasting on Active Virtual Private Networks," *Proceedings of the 25th Annual IEEE Conference on Local Computer Networks (LCN 2000)*, pp. 213-222.
- [LeSt98] W. Lee, S.J. Stolfo, "Data Mining Approaches for Intrusion Detection," *Proceedings of the 7th USENIX Security Symposium (SECURITY 1998)*, San Antonio, TX, Jan. 1998.
- [Level3] Level 3 home page, <http://www.level3.com>.

- [LHKJ00] H. Lee, J.Hwang, B. Kang, K. Jun, "End-to-end QoS architecture for VPNs: MPLS VPN deployment in a backbone network," *Proceedings of the 2000 International Workshop on Parallel Processing*, 2000, pp. 479-483.
- [LIDS] Xie Huagang, "Build a Secure System with LIDS," Linux Intrusion Detection System (LIDS) home page, http://www.lids.org/document/build_lids-0.2-1.html (updated October 2000).
- [LiPo98] U. Lindqvist, P.A. Porras, "Detecting Computer and Network Misuse through the Production-Based Expert System Toolset (P-BEST)," *Proceedings of the IEEE Symposium on Security and Privacy*, Oakland, CA, May 1999, pp. 146-161.
- [Loal98] P.A. Loscocco, et al., "The Inevitability of Failure: The Flawed Assumption of Security in Modern Computing Environments," *Proceedings of the 21st National Information Systems Security Conference*, October 1998, pp. 303-314.
- [LPS99] W. Lee, C.T. Park, S.J. Stolfo, "Automated Intrusion Detection Methods Using NFR," *Proceedings of the 1st USENIX Workshop on Intrusion Detection and Network Monitoring*, April 1999.
- [LSC97] P.K. Chan, "Learning Patterns from Unix Process Execution Traces for Intrusion Detection," *AAAI Workshop: AI Approaches to Fraud Detection and Risk Management*, AAAI Press, July 1997.
- [LSC01] W. Lee, S.J. Stolfo, P.K. Chan, E. Eskin, W. Fan, M. Miller, S. Hershkop, J. Zhang, "Real Time Data Mining-based Intrusion Detection," *Proceedings of DISCEX II*, June 2001.
- [LSM98] W. Lee, S.J. Stolfo, K.W. Mok, "Mining Audit Data to Build Intrusion Detection Models," *Proceedings of the 4th International Conference on Knowledge Discovery and Data Mining (KDD 1998)*, AAAI Press, Aug. 1998.
- [LSM99] W. Lee, S.J. Stolfo, K.W. Mok, "A Data Mining Framework for Building Intrusion Detection Models," *Proceedings of the 5th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (ICDD 1999)*, San Diego, CA, Aug. 1999.
- [MaYu99] A. Mayer, M. Yung, "Secure protocol transformation via 'expansion': from two-party to groups," *Proceedings of the 6th ACM Conference on Computer and Communications Security*, Singapore, November 1999, pp. 83-92.
- [Metz99] C. Metz, "AAA Protocols: Authentication, Authorization, and Accounting for the Internet," *IEEE Internet Computing*, Nov./Dec. 1999, pp. 75-79, <http://computer.org/internet>.
- [MiLuBr97] T. Mitchem, R. Lu, R.D. Brien, "Using Kernel Hypervisors to Secure Applications," *Proceedings of ACSAC '97*, <http://www.securecomputing.com/khyper/acsac97.pdf>.

- [MIT99] Information Systems Technology Group, MIT Lincoln Laboratory, "DARPA Intrusion Detection Evaluation," 1999, <http://www.ll.mit.edu/IST/ideval/index.html>.
- [Mosko97] R. Moskowitz, "The Complicated World of Digital Signatures," Tech Web home page, <http://www.networkcomputing.com/819/819colmoskowitz.html> (October 1997).
- [MPC98] R. Meo, G. Psaila, S. Ceri, "An Extension to SQL for Mining Association Rules," *Data Mining and Knowledge Discovery 2*, Kluwer Academic Press, Boston, 1998, pp. 195-224.
- [MTV97] H. Mannila, H. Toivonen, A.I. Verkamo, "Discovery of Frequent Episodes in Event Sequences," *Data Mining and Knowledge Discovery 1*, Kluwer Academic Publishers, 1997, pp. 259-289.
- [MWV00] W.B. Martin, P.D. White, W.M. Vanfleet, "Government, Industry and Academia: Teaming to Design High Confidence Information Security Applications," *FMSP '00*, Portland, Oregon, pp. 37-46.
- [NeFi] C.L. Nelson, D.S. Fitzgerald, "Sensor Fusion for Intelligent Alarm Analysis," 0-7803-3537-6-9/96, IEEE Press, 1996, pp. 143-150.
- [Net2K2NE] NetK2NE home page, "Leased Line Startup Costs," <http://www.k2nesoft.com/leased.html>.
- [Neum00] P.G. Neumann, "Practical Architectures for Survivable Systems and Networks," Report on U.S. Army Research Laboratory (ARL), Contract DAKF11-97-C-0020, June 2000, <http://www.csl.sri.com/users/neumann/arl-one.html>.
- [NIAP] National Information Assurance Partnership (NIAP) home page, <http://niap.nist.gov/>.
- [Openwall] Openwall project, "Security Patch for Linux Kernel," <http://www.openwall.com/linux>.
- [Opli96] R. Oplinger, *Authentication Systems for Secure Networks*, Artech House Computer Security Series, 2000.
- [Opli98] R. Oplinger, *Internet and Intranet Security*, Artech House Computer Security Series, 2000.
- [Opli00] R. Oplinger, *Security Technologies for the World Wide Web*, Artech House Computer Security Series, 2000.
- [Owl] Openwall GNU/*/Linux (Owl) home page, <http://openwall.com/Owl/>.
- [PASW98] J.L. Abadpeiro, N. Asokan, M. Steiner, M. Waidner, "Designing a Generic Payment Service," *IBM Systems Journal*, vol. 37, no. 1, 1998, pp. 72-88.

- [PDM] S. Farrell, R. Perlman, C. Kaufman, M. Rose, "Securely Available Credentials - The PDM Protocol," IETF Securely Available Credentials Working group, Internet draft draft-ietf-sacred-protocol-beep-pdm-00.txt, June 2001 (expires: January 2002), <http://www.ietf.org/internet-drafts/draft-ietf-sacred-protocol-beep-pdm-00.txt>.
- [PKCS15] RSA Laboratories, "PKCS#15 Cryptographic Token Information Format Standard," <http://www.encryption.com/rsalabs/pkcs/pkcs-15/>.
- [Penn00] D. Penn, "Secure Computing to Develop Type Enforced Tux," *Linux Journal* home page, <http://www2.linuxjournal.com/articles/briefs/059.html> (February 2000); <http://www.linuxjournal.com/article.php?sid=5105>.
- [Perry98] G. Perry, email message to "The Intelligent hacker's Choice! Firewall Archives," Posted Jan. 1998, <http://www.netsys.com/firewalls/firewalls-9801/0577.html>.
- [PKIX] Internet Engineering Task Force, "Public Key Infrastructure X.509," <http://www.ietf.org/html.charters/pkix-charter.html>.
- [Princ01] Princeton Software home page, "What is Active Data Archiving?" http://www.storesmarter.com/b0_whatitis.htm (September 2001).
- [Privs] Linux Kernel Archives, "Linux-Privs," <http://www.kernel.org>.
- [RFC2078] J. Linn, "Generic Security Service Application Program Interface (GSS API)," RFC2078, January 1997, <http://www.cis.ohio-state.edu/cgi-bin/rfc/rfc2078.html>.
- [RFC2138] C. Rigney, A. Rubens, W. Simpson, S. Willens, "Remote Authentication Dial In User Service (RADIUS)," RFC 2138, <http://www.cis.ohio-state.edu/cgi-bin/rfc/rfc2138.html> (April 1997).
- [RFC2487] P. Hoffman, "SMTP Service Extension for Secure SMTP over TLS," RFC2487, January 1999, <http://www.cis.ohio-state.edu/cgi-bin/rfc/rfc2487.html>.
- [RFC2945] T. Wu, "SRP Authentication and Key Exchange System," RFC2495, September 2000, <http://www.ietf.org/rfc/rfc2945.txt>.
- [Rick97] J. Rickard, "Internet Architecture," *Boardwatch*, 1997, <http://www.pslash.org/~vool/articles/bw-netarchitecture-1997.html>.
- [SafeStart] Integrity Sciences, Inc., "SafeStart: A Fail-safe Integrity Checking System," last updated 1998, <http://world.std.com/~dpi/safestar.html>.
- [SDSI] Internet Engineering Task Force, "Simple Distributed Security Infrastructure," http://www.ietf.org/html.charters/draft-ietf-spki-cert-theory-*.txt.

- [SecPortal] Linux Documentation Project, Linux Administrator Security Guide, Patches for Linux Kernel, <http://securityportal.com/lasg/kernel>, <http://www.linuxdoc.org/LDP/lasg/kernel>.
- [Secsh] Secure Shell (secsh) IETF Charter, SSH Working group home page, <http://www.ietf.org/html.charters/secsh-charter.html> (updated July 2001).
- [SAS] SAS Institute home page, <http://www.sas.com>.
- [SCVPN] Secure Computing, "White Paper; an Overview of Virtual Private Networks (VPNs)," <http://www.securecomputing.com/index.cfm?sKey=444> (March 2000).
- [SELinux] The National Security Agency, "Security-Enhanced Linux," <http://www.nsa.gov/selinux/index.html>.
- [SGGB99] E.G. Sirer, R. Grimm, A.J. Gregory, B.N. Bershad, "Design and Implementation of a Distributed Virtual Machine for Networked Computers," SOSP-17 12/1999, Kiawah Island, SC, *ACM 1-58113-140-2/99/0012*, pp. 202-216.
- [SIAR00] WetStone Technologies, Inc., "SIAR-IA: Strategic Intrusion Assessment Research and Integration Architecture," Final Report, Contract #F30602-00-C-0153, Air Force Research Laboratory, Rome, NY, July 2000.
- [Sidewn] Sidewinder team, "Secure Computing's e-Security Challenge Results," Secure Computing Corporation Web page, <http://www.securecomputing.com/index.cfm?sKey=803> (updated July 2001).
- [Signal98] *Signal* magazine, "Programmable Cryptography Emerges from Advanced Chip," August 1998, <http://www.us.net/signal/Archive/Aug98/program-aug.html>.
- [sLinux] sLinux project, The Secure Internet OS (sLinux), <http://www.slinux.com>.
- [SMN99] Portland State University, "Secure Mobile Networking Project," <http://www.cs.pdx.edu/research/SMN>.
- [SOCKS] M. Leech et al., "SOCKS version 5 Protocol," RFC1928, March 1996, <http://www.socks.nec.com/rfc/rfc1928.txt>.
- [SOCKSF] NEC Networking Systems, "SOCKS FAQs," <http://www.socks.nec.com/socksfaq.html> (last updated September 2000).
- [SPREAD] Y. Amir, J. Stanton, "The Spread Wide Area Group Communication System," Department of Computer Science, The Johns Hopkins University, Technical Report CNDS 98-4, <http://citeseer.nj.nec.com/84858.html>.

- [SmRh00] J.F. Smith III, R.D. Rhyne II, "Genetic algorithm based optimization of a fuzzy logic resource manager for electronic attack," *Proceedings of SPIE vol. 4057 (2000) In Data Mining and Knowledge Discovery: Theory, Tools, and Technology 2*, 2000, 0277-786X/00, pp. 62-71.
- [Snort01] Open Source Network Intrusion Detection System (Snort) home page, <http://snort.sourceforge.com/>.
- [Sokol99] L. Sokol, "Data Mining in the Real World," *Part of the SPIE Conference on Data Mining and Knowledge Discovery: Theory, Tools, and Technology*, 1999, SPIE vol. 3695, 0277-786X/99, pp. 192-196.
- [SPW98] L. Sacks et al., "TRUMPET Service management Architecture," *Proceedings of the 2nd International Enterprise Distributed Object Computing Workshop, EDOC '98*, 1998, pp. 289-295.
- [SRP] T. Wu, "The Stanford SRP Authentication Project," Updated August 2001, <http://www-cs-students.stanford.edu/~tjw/srp/analysis.html>.
- [SSH] Secure Shell (SSH), <http://www.ssh.fi>.
- [Ste97] L. Stein, *Web Security*, Addison Wesley Longman Inc., Reading, MA, 1997.
- [SVS] SVS Internet Services home page, <http://www.svs.com/svs-info/business/lan-56.html>
- [SYJS00] D. Shands, R. Yee, J. Jacobs, E.J. Sebes, "Secure Virtual Enclaves: Supporting Coalition Use of Distributed Application Technologies," *Proceedings of the Network and Distributed Security Symposium (NDSS 2000)*, San Diego, CA, 2000.
- [TeleWeb] TeleWeb home page, <http://www.teleweb.net/twplans.htm>.
- [Ther01] Spyrus products page, "Terisa TLS Gold," <http://www.spyrus.com/content/products/Terisa/TLSGold.asp>.
- [Touc00] J. Touch, "Dynamic Internet overlay Deployment and Management Using the X-Bone," *Proceedings of the International Conference on Network Protocols*, 2000, pp. 59-68.
- [TSEC] U.S. Department of Defense, *Trusted Network Interpretation of the Trusted Computer System Evaluation Criteria*, Fort George G. Meade, MD, 1987.
- [Trustix] Trustix Secure Linux, The Server Distribution, <http://www.trustix.net> (updated April 2001).
- [UFN] The Userfriendly Network, "Secure Patch for Linux Kernel," <http://niteowl.userfriendly.net/linux>.

- [WaSc96] D. Wagner, B. Schneier, "Analysis of the SSL 3.0 protocol," *Proceedings of 2nd USENIX Workshop on Electronic Commerce*, USENIX press, November 1996, pp. 29-40.
- [Wets00] WetStone Technologies, Inc., "SIAR-IA: Strategic Intrusion Assessment Research and Integration architecture Strawman Architecture proposal," Report prepared for Air Force Research Laboratory/IFGB, Contract #F30602-00-C-0153, July 28, 2000.
- [Wets01] WetStone Technologies, Inc., "NET-FLARE: Network Fuzzy Logic Attack Recognition Engine," Manual prepared for Air Force Research Laboratory/IFGB, Contract #F30602-00-C-0044, June 30, 2001.
- [Willow] Software Engineering Research laboratory, University of Colorado, "Tolerating Intrusions Through Secure System Reconfiguration," <http://www.cs.colorado.edu/serl/its> (last updated 2000).
- [WTG00] W. Martin, P. White, F.S. Taylor, A. Goldberg, "Formal construction of the Mathematically Analyzed Separation Kernel," *Proceedings of the 15th IEEE International Conference on Automated Software Engineering*, Sept. 2000, pp. 133-141.
- [XFN] Open Group, "Federated Naming: The XFN Specification," Open Group Technical Standard C403 ISBN 1-85912-052-0, July 1995.
- [Youn00] R. Younglove, "Virtual Private Networks: How They Work," *Computing and Control Engineering Journal*, December 2000, pp. 260-262.
- [Young00] R. Younglove, "Virtual Private Networks: Secure Access for E-Business," *IEEE Internet Computing*, July/Aug. 2000, pp. 96-99.
- [ZDNet] ZDNet home page, "Will You Pay Internet Tolls?" <http://www.zdnet.com/sp/stories/news/0,4538,2341710,00.html>.
- [ZhLa98] J. Zhou, K-Y. Lam, "Undeniable Billing in Mobile Communications," *MOBICOM'98*, Dallas, TX, 1998, ACM 1998 1-58113-035-x/98/10, pp. 284-290.