MP 99B0000020

MITRE PRODUCT

State of the Art in Anomaly Detection and Reaction

July 1999

Leonard J. LaPadula

Sponsor: Air Force
Dept. No.: G021
Contract No.: F19628-99-C-0001
Project No.: 03997482

MITRE
Center for Integrated Intelligence Systems
Bedford, Massachusetts

Abstract

This paper presents a view of the state of the art in anomaly detection and reaction technology. The paper develops the view from six sources: three prior reports (two national, one MITRE), a survey of commercially available software, a survey of government software, and a survey of government-funded research projects. The author performed the surveys for this paper.

KEYWORDS: anomaly detection and reaction, state of the art, intrusion detection and reaction, vulnerability scanning, policy compliance scanning, network monitoring, host monitoring, anomaly response, intrusion response

Table of Contents

Section	Page
State of the Art in Anomaly Detection and Reaction	**1**
Introduction	1
About Anomaly Detection and Reaction Systems	1
About Describing the State of the Art	2
How This Paper is Organized	4
National Info-Sec Technical Baseline	4
Report of Hill and Aguirre	6
Intrusion Detection Subgroup's Report	6
Commercial Products	7
Summarized Information About the Tools in the Compendium	8
Comparison of Tools to Issues Identified in Earlier Reports	10
Government Products	14
Research Efforts	15
Summary	18
Methods of Detection	25
Statistical Deviation Detection	25
Pattern Matching	26
Use of the Methods	27
Types of Tools	27
Architectural Attribute of Tools	30
List of References	**23**
Appendix: About Anomaly Detection and Reaction Systems	**25**
Distribution List	**31**

State of the Art in Anomaly Detection and Reaction

Introduction

This paper synopsizes the state of the art in anomaly detection and reaction systems. Having said that necessitates some explanation about what we mean by "anomaly detection and reaction" and what we mean by "state of the art."

About Anomaly Detection and Reaction Systems

We are interested in automated capabilities that can detect or find anomalies in computer systems, report them in useful ways, remove discovered anomalies, and repair damage they may have caused. Included in this scope of interest are traditional intrusion detection and reaction tools. The broader scope of anomaly detection and reaction also includes vulnerability scanners, infraction scanners, and security compliance monitors. These tools protect not only against intruders but against errors and carelessness in administration and operation of end systems and network components.

Even within the narrower scope of intrusion prevention, one can fashion protection against cyber intruders within a spectrum of techniques. At one end of the spectrum is the method of detecting intruders. In this method, one uses intrusion detection tools to watch what is going on in the network to discover suspicious events. If perfect intrusion detection and reaction systems were available, there might be no need for any other measures to protect against cyberattack. At the other end of the spectrum is the method of ensuring that all the components of the network, including firewalls, routers, servers, and workstations, are equipped to fully repel any attack. In this method, one does not try to detect intrusive connections coming from outside one's network since they can do no harm. In theory, even a denial of service attack can be thwarted in this way because the components of the data communications infrastructure would be smart enough not to carry the traffic that would cause the denial of service. Of course, it is a good question to ask how to make the components so smart. In practice, neither end of the spectrum will provide the best protection for investment made. Prudence demands balance: organizations should properly set up and configure the components of their networks using current best practices and they should provide state of the art intrusion detection capability, depending wholly on neither to adequately protect their computing resources. The capabilities brought to bear by the nondetection tools protect not only against intruders but against errors and carelessness in administration and operation of both end systems and network components.

Doing these things is not a one-time chore. Network topologies tend to be dynamic. Often it is difficult to control the comings and goings of hosts on a network, especially in large networks. The job of properly setting up and configuring components often requires skilled

personnel, who are in short supply. In addition, new cyberattacks may demand new protections or responses.

Prudent, affordable, continuous protection of one's network involves monitoring the network for anomalies of various kinds, whether they are suspicious textual strings in a network packet or undesirable values for important keys in NT registries. Moreover, it involves correcting detected anomalies, whether that means terminating a connection or reconfiguring a server.

We call an automated system that performs or assists in such tasks an anomaly detection and reaction (ADR) system. Besides checking network packets for suspicious strings, or monitoring a user's behavior looking for deviations from an established pattern, the ADR system checks components of the network for errors of omission, misconfigured applications, and errors in system parameters. When the ADR system finds an anomaly, it reacts, generally by trying to fix the anomaly. Its response may be restricted to issuing an alert for certain anomalies. For others, it may be able to fully correct the problem. In some cases, it may be able to provide ancillary information that will assist an administrator in correcting the anomaly. What it can do will be determined by the state of the art, the budget, and the information operation to be protected.

Besides budgetary considerations, the extent of the protection domain determines the needed capacity of the ADR system for that domain. Networks tend to grow, thereby extending the scope of interest for an ADR system. Thus, scalable ADR systems are needed, not only so that the same basic system can serve domains of different size, but also so that it can accommodate significant growth in the domain it protects.

The appendix in this report contains a discussion of related ideas on anomaly detection and reaction, describing methods of anomaly detection and providing a classification of ADR tools.

About Describing the State of the Art

A description of the state of the art in some technology generally includes more than what has been reduced to ordinary practice. State-of-the-art reports may use research efforts, for example, to describe the level of sophistication achieved in some developing technology, such as the technology that deals with nanocomputers. A danger is that one may take as "reduced to practice" what is merely a possibility. We too will include more than "practice" in this report, but we will focus on what is commercially available today in ADR. Besides wanting to avoid confusing "practice" with "theory", we are motivated by a concern for the particular needs of our military customers. We want a pragmatic view that is relevant to the Air Force's mission-oriented networks and computers, for example. This view will enable us to tell what is currently available that can provide an information-assurance benefit while fitting within the constraints of that environment.

We have drawn on three reports dealing with state of the art, which we will identify shortly. These reports deal with intrusion detection and reaction, not the broader anomaly detection and reaction of interest here. Thus, they do not consider vulnerability scanners, infraction scanners, and security compliance monitors.

We present a broader picture in this synopsis and bring it up to date by summarizing what we have learned from current[1] descriptions of commercial off-the-shelf (COTS) products. This part of our synopsis provides the pragmatic view we talked about earlier. Government off-the-shelf (GOTS) products also help define the current state of the art because they may provide capabilities not yet found in commercial products or, at least, their use may shed light on the economics affecting their users. In addition to a summary of GOTS products, we include a look ahead by summarizing what we know of current research efforts funded by the U.S. government. Note that this paper does not include information about freeware, shareware, research by vendors, ad hoc consortia or teams that may be implementing ADR tools, or any other source not explicitly identified herein. Having so noted, readers should understand the exact scope of the state of the art as reported here and be able to judge the applicability of this synopsis for each situation they may deal with.

In short, this synopsis draws on these sources of information

- The National Info-Sec Technical Baseline report on intrusion detection and response [1]
- The description of the state of the art in network-based intrusion detection systems in a report of Hill and Aguirre [2]
- The report of the Intrusion Detection Subgroup of the National Security Telecommunications Advisory Committee on the implications of intrusion detection technology research and development on national security and emergency preparedness [3]
- Product descriptions of commercial off-the-shelf (COTS) and government off-the-shelf (GOTS) ADR systems [4]
- Descriptions of current research in anomaly detection and reaction [4]

---

1 As of March, 1999

How This Paper is Organized

We begin by summarizing the conclusions from the three reports, in chronological order. Then we take a look at commercial products, government products, and research efforts. Finally we capture the state of the art as a summary of the preceding information sources. Thus, the rest of this paper is organized as follows:

- National Info-Sec Technical Baseline: summary of findings
- Report of Hill and Aguirre: summary of findings
- Intrusion Detection Subgroup's Report: summary of findings
- Commercial Products: summary of product types and characteristics
- Government Products: summary of product types and characteristics
- Research Efforts: summary of principal lines of investigation
- Summary: a capsule description of the state of the art in anomaly detection and reaction

National Info-Sec Technical Baseline

In the Executive Summary of the baseline report, the 1996 state of the art in intrusion detection is succinctly characterized in the first paragraph

> "The state of the art in logical intrusion detection of national information infrastructure (NII) systems is such that a human expert working with a well-developed set of tools can implement a detection system in a few months for a special-purpose computing environment with the properties that it
>
> (1) reliably detects a substantial number of known intrusion techniques,
>
> (2) detects substantial short-term changes in user and system behavior,
>
> (3) produces many alarms that, on investigation, are not intrusions (false-positives), and
>
> (4) fails to alarm on an unknown number of intrusions (false-negatives)."

Similarly, the report summarizes the situation in automated response as follows:

> "The state of the art in automated response to detected intrusions is that we can program a wide variety of responses, ranging from increased defenses to offensive counter-strikes. Examples throughout this range have been demonstrated. Several important automated response issues remain essentially unaddressed at this time, including, but not limited to, (1) limiting the effect of automated response so as to prevent cascade and livelock failures that may be caused by the response system, (2) providing safeguards against false-positives and enemy-induced erroneous responses, and (3) using the response system to push the point of attack deflection back toward the attack source."

The report identifies these practical issues

- Testing is lacking, with most systems rarely tested beyond demonstrating that they detect some anomalous events
- Damage assessment and recovery mechanisms are lacking in the vast majority of systems
- Scalability of IDR systems is severely limited. To be of significant utility in modern information environments, IDR systems must be capable of handling large numbers of events. By contrast, most of the systems examined by the reporters detect events at the system level. Work directed toward correlating activities between systems is in the very early stages and much more work needs to be done.

The report stresses the importance of identifying the source of an attack and states that traceback capability is most difficult in the computer-networking environment.

> "In today's Internet, there are no central controls or methods to trace through the infrastructure without the cooperation of a potentially large number of independent infrastructure providers. In recent years, some limited tools have been developed to try to trace an intrusion to a source, but these tools are only effective against the least sophisticated attackers. They do not allow traceback when IP address forgery is used, when intermediate nodes fail to cooperate[2], when firewalls block further traceback, or when the intruder breaks into an intermediate site in order to launch the attack."[3]

---

[2] We note that traceback techniques that do not necessarily require the cooperation of the independent infrastructure providers have a much better chance, technically speaking, to locate the cybersource of an attack; the United States Justice Department, however, typically does not allow military departments to use such techniques.

[3] The report cites Cohen's paper *A Note on Distributed Coordinated Attacks*, Computers and Security, August 1996.

Report of Hill and Aguirre

The state of the art for network-based intrusion detection was described in September 1997 by Hill and Aguirre. [7] In summary, they reported that[4], for network-based intrusion detection,

- No common terminology or taxonomy exists for discussing intrusion detection techniques or implementation
- Current network-based intrusion detection systems (IDSs) are signature recognition systems
- The IDSs can be modeled as being performed by three agents
  - Collection function
  - Analysis function
  - Management function
- The collection function needs improvement: it should address new physical media, such as ATM and Fast Ethernet, provide more flexibility in signature specification, and understand more application protocols.
- Better analysis support is even more urgent a need:
  - Improved expert systems should reduce the false alarm rate
  - Better distribution of problem identification is needed
  - Tools that can correlate activity from distributed sources and activity directed toward distributed targets are crucial to address issues of information warfare.
- There is growing recognition that there would be high utility in integrating the output of different entities involved in network security, including routers, firewalls, proxies, and host-based and network-based IDSs.

Intrusion Detection Subgroup's Report

This study focused on influencing research and development of intrusion detection technology so that national security and emergency preparedness (NS/EP) requirements can be satisfied. The findings of the subgroup on technology relevant to the telecommunications infrastructure were

- Research and development seldom focuses on controlling elements of the telecommunications infrastructure; IDSs in use mostly were developed to meet the needs of host systems in unique environments (for example, UNIX, Novell, Microsoft NT)

---

[4] As of September 1997.

- "Many of the deployed IDSs are unable to scale to the large environments characterized by thousands of network nodes, which limits their ability to detect intrusions effectively across different platforms, applications, networks, and infrastructures."
- Standards, testing, and validation procedures are needed to enable verification of IDSs' capabilities; the development and use of standard metrics would be helpful in this area
- False alarms are a major performance problem in current IDSs
- The fact that current IDSs can detect only what they have been preprogrammed to detect, like virus detectors, limits their effectiveness in the face of new attacks
- Detection and response need to occur in real time to limit potential damage; current IDSs fall far short of this ideal
- No systems appeared to provide an automated damage assessment and response capability

Commercial Products

We base our findings in this section on product information gathered primarily from vendors' web sites, as reported in the ADR Compendium [4]. The ADR Compendium includes information about 39 commercially available products from a number of vendors. In the compendium, vendors have been grouped as primary providers, secondary providers, and others. Primary providers are those with the highest revenues as reported in the Hurwitz Group white paper *Information Security: Assessing Risks and Detecting Intrusions.* Secondary providers are those with comparable, competitive tools or systems, as identified in the same paper. Other vendors offering ADR tools have also been included.

The primary providers are

- AXENT Technologies
- CISCO (recently acquired WheelGroup)
- Internet Security Systems (ISS)
- Intrusion Detection, Inc., a Security Dynamics Company
- Network Associates, Inc. (recently acquired Trusted Information Systems (TIS))
- PLATINUM Technology

As is suggested from the list of names, the primary providers tend to be large, well-established vendors who have either been in the intrusion detection area for a while or have

acquired significant capability through acquisition. These same vendors tend to have the most recent products or, in some cases, recently upgraded versions of older products.

Summarized Information About the Tools in the Compendium

The next three tables summarize selected information from the compendium. Table 1 shows counts of the number of products organized by type of product and category of vendor (P: Primary; S: Secondary; O: Other; as explained above). Types of products are further subdivided, where appropriate, into architectural subtypes. These subtypes are explained in the appendix of this report, Architectural Attribute of Tools.

Table 1. Summary of Types of Tools by Vendor Categories

Product Type	Architecture	P	S	O
ADR Director	Director	1		
Anomaly Detection Support Tool	Sensor			2
Network Mapper + Vulnerability Scanner + Risk Analyst	Sensor		1	
Network Monitor	Sensor			4
	Agents-Director	3		
	unknown			1
Network Monitor + Infraction Scanner	Agents-Director	1		
Security Compliance Scanner	Sensor	1		1
	Agent			1
	Director			1
	Agents-Director	1		
Suite of Monitors (web access, LAN packets, tracing)	Sensor			1
System Monitor	Sensor	2		2
	Agents-Director	3		1
	Sensors-Director	1	1	
System Monitor + Vulnerability Scanner	Agents-Director		1	
System Monitor for Access Control	Sensors-Director	1		
Vulnerability Scanner	Sensor	5	1	
Vulnerability Scanner with built-in Infraction Scanning	Agents-Director	1		

Vulnerability Scanner with built-in Network Mapping	Sensor	2		

The tools listed in this table go well beyond intrusion detection, with functionality extending into areas of analysis that just several years ago were untouched by the commercial vendors. Table 2 shows the major functions of the ADR products reported in Table 1 with a count for each function of how many products provide that function. In this table, a product may be counted more than once. If, for example, it provides both vulnerability scanning and network mapping, it counts in both categories.

Table 2. Counts of Products Providing the Major Functions

Function	P	S	O
ADR Management	1		
ADR Support			2
Network Mapping	2	1	
Risk Analysis		1	
Network Monitoring	4		6
Infraction Scanning	2		
Security Compliance Scanning	2		2
System Monitoring	7	2	3
Vulnerability Scanning	8	2	

Table 2 shows evidence of the broadening functionality of the tools. Recall our comment earlier that the primary providers generally have the more recent tool offerings[5]. Taking the Primary, Secondary, Other column headings in the table as a rough time line extending backward in time, we observe a marked increase in the number and diversity of functions provided. The vulnerability scanning function stands out in this regard. Older tools primarily provided network and system monitoring capability, the core functions of intrusion detection. We see now a burgeoning supply of anomaly detection capabilities.

---

[5] One of the Security Compliance Scanning tools listed under the Other column is, in fact, a new tool recently made available (i.e., Microsoft's Security Compliance Monitor).

Table 3 displays the number of products in the three major architectural types, organized by provider. The three architectural types are Sensor/Agent (includes Sensor and Agent types in Table 1), Director, and Sensors/Agents-Director (includes Sensors-Director and Agents-Director types in Table 1).

Table 3. Numbers of Products by Architectural Types

Architecture	P	S	O
Sensor/Agent	11	2	11
Director	1		1
Sensors/Agents-Director	11	2	1

The data of Table 3 indicate an increasing use of the Agent/Sensor-Director architecture in more recent products. This certainly agrees with one's intuitive notion of how such products might evolve. With the use of the Agent/Sensor-Director architecture comes improved scalability of the products in some cases. For example, when a Director unit can accept reports and inputs from other Director units, the possibility exists for a hierarchical arrangement of ADR capability with wide coverage at the "leaf" level.

Comparison of Tools to Issues Identified in Earlier Reports

We can get a good sense of how the state of the art is evolving by examining issues identified earlier in light of today's capabilities. In Table 4 we list each of the issues identified in the three reports summarized earlier in this paper. The source of each issue is indicated by a short reference as follows:

- NISTB: the National Info-Sec Technical Baseline report
- H&A: the report of Hill and Aguirre
- IDSR: the Intrusion Detection Subgroup's Report

Beside each issue, we provide commentary based on our understanding of available COTS tools.

Table 4. Comparison of Tools to Issues Identified in Earlier Reports

Issue	Commentary
(NISTB) Detection of known intrusion techniques	The NISTB report indicated that the tools could reliably detect a substantial number of known intrusion techniques. Today that number is in the range of four to five hundred known patterns for pattern-matching network monitors.
(NISTB) False positives (IDSR) False alarms are a major performance problem in current IDSs	We have seen no evidence of significant improvement in avoiding false positives at the detection level; this problem may be ameliorated by improved capability for user customization and analysis. The performance may now be more of an analysis problem than a capture problem.
(NISTB) False negatives	There are known cases of false negatives occurring with esoteric attacks. The situation has improved, however, for run-of-the-mill attacks as the tools increasingly provide comprehensive coverage of known techniques and increasingly more frequent and user-friendly pattern updates.
(NISTB) Limiting the effect of automated responses to prevent tool-induced failures; providing safeguards in the automated response system against false positives and enemy-induced erroneous responses; and using the response system to push the point of attack deflection back toward the attack source.	The NISTB report indicated that these areas were essentially unaddressed at that time (December 1997). Current tools show no direct evidence of progress in these areas. However, they generally provide user interfaces that allow some customization, which may make it easier for an administrator or operator to have better control over these factors.

Issue	Commentary
(NISTB) Testing is lacking (IDSR) Standards, testing, and validation procedures are needed to enable verification of IDSs' capabilities	There are no testing tools in the arsenal of COTS ADR tools, neither built-in nor as part of a suite, with one exception among the products examined in the ADR Compendium [4]. The vendor of one recent vulnerability scanning product provides software for setting up a fake DNS server to enable the scanning product to check for the DNS server cache-overflow vulnerability; the scanner also comes with a scripting language that allows users to create specialized network packets for vulnerability testing.
(NISTB) Damage assessment and recovery mechanisms are lacking in the vast majority of systems (IDSR) No systems appeared to provide an automated damage assessment and response capability	There has been little change in this area; the COTS tools we have studied generally provide no damage assessment and recovery capabilities. Two products provide some capability for automatic repair of illicit changes. One uses a fairly general technique that can interface to a variety of network management systems.
(NISTB) Scalability of IDR systems is severely limited (IDSR) Many of the deployed IDSs are unable to scale to the large environments characterized by thousands of network nodes, which limits their ability to detect intrusions effectively across different platforms, applications, networks, and infrastructures.	Scalability has improved, especially because of the increasing use of the Sensors/Agents-Director architecture. Three vendors explicitly address this area, providing security compliance scanning, system monitoring, and network monitoring capabilities that scale up to enterprise-wide networks (e.g., WANs). The scale of operation among the tools available today is improved over what it is was two or three years ago; some of today's intrusion detection tools can cover large environments of thousands of nodes. As these tools come into use, the situation noted in IDSR should gradually change.

Issue	Commentary
(NISTB) Traceback capability	None of the intrusion detection tools examined provided any capability in this area.
(H&A) No common terminology or taxonomy exists for discussing intrusion detection techniques or implementation	The situation has improved; the vendors of the COTS tools examined in this study are essentially speaking the same language in describing their offerings.
(H&A) The collection function in intrusion detection systems needs improvement: it should address new physical media such as ATM and Fast Ethernet and provide more flexibility in signature specification. (IDSR) Current IDSs can detect only what they have been preprogrammed to detect, like virus detectors; this limits their effectiveness in the face of new attacks.	The intrusion detection tools examined in this study generally cover Ethernet, Fast Ethernet, and FDDI network topologies; none directly address ATM. Some tools allow the user to add signature specifications or rules. In addition, many vendors are now providing updates on a regular basis (at least monthly); the trend is toward automatic updates via the Internet as new attacks are codified; typically today there is an e-mail notification and updates can be downloaded manually.
(H&A) Better analysis support is needed: improved expert systems should reduce the false alarm rate; and tools that can correlate activity from distributed sources and activity directed toward distributed targets are crucial to address issues of information warfare.	Some progress has been made: a number of tools are able to collect data from distributed sources for presentation to the user; some correlation capability, to correlate multiple attacks and to correlate attacks with vulnerabilities, is now available.
(H&A) There is growing recognition that there would be high utility in integrating the output of different entities involved in network security, including routers, firewalls, proxies, and host-based and network-based IDSs.	A discernible trend in this direction has developed among the COTS products. Those products that are bundled in suites typically can integrate the output of multiple, different sensors at a manager station.

Issue	Commentary
(IDSR) Detection and response need to occur in real time to limit potential damage; current IDSs fall far short of this ideal	Although responses in real time are limited in nature, they can occur rapidly enough to avoid or limit damage (e.g. shutting down a session, terminating a process, shunning a connection).

Government Products

We base our findings in this section on product information gathered primarily from providers' web sites, as reported in the ADR Compendium [4]. The ADR Compendium includes information about three products available from

- Air Force Information Warfare Center (AFIWC/AFCERT), providing Automated Security Incident Measurement (ASIM), Version 2.0
- Lawrence Livermore National Laboratory, providing Network Intrusion Detector (NID), Version 2.1
- DISA Information Assurance Support Environment (IASE), providing Joint Intrusion Detection System (JIDS), Version 2.0.3

JIDS is the DoD version of NID, providing essentially the same functionality as NID. These tools are considered network monitors by their providers. In the terminology of this report, ASIM running in batch mode is considered an infraction scanner, in real-time mode a network monitor. NID (and JIDS) are architecturally sensors; ASIM uses an agents-director architecture.

These tools are very similar in functionality, essentially performing an intrusion detection function by analyzing Ethernet and FDDI packets. The major difference is that NID (and JIDS) sensors deal with a single security domain while the ASIM director unit deals with multiple security domains, each of which has an ASIM sensor that reports to the single director unit.

What do we learn about the state of the art from these tools? The fact that they came into being in the first place indicates that a need existed that could not be met by commercial products within the budgetary constraints in place at that time. This appears to continue to be a factor, especially with regard to the scale on which ASIM operates. Aside from this consideration, the GOTS products appear to offer no capability that is technically superior to that of the COTS products and, in fact, appear now to be lagging in providing some significant features, such as SNMP traps for alerting, automatic updating of signatures, and user-customization capabilities.

Research Efforts

In this section we use information gathered from researchers' web sites, as reported in the ADR Compendium [4]. The research efforts reported in the Compendium are

- Automated Intrusion Detection Environment (AIDE) Advanced Concept Technology Demonstration (ACTD)
- Autonomous Agents for Intrusion Detection (AAFID)
- Common Intrusion Detection Framework (CIDF)
- DARPA Intrusion Detection Evaluation
- Distributed Intrusion Detection System (DIDS)
- Event Monitoring Enabling Responses to Anomalous Live Disturbances (EMERALD)
- Extensible Prototype for Information Command and Control (EPIC^2)
- Graph-based Intrusion Detection System (GrIDS)
- Next-Generation Intrusion Detection Expert System (NIDES)
- Spitfire

In Table 5, we briefly state the main thrust of each effort and comment on it with respect to how it participates in defining the state of the art.

Table 5. Main Thrusts of Research Efforts

Project	Thrust	Comment
Automated Intrusion Detection Environment (AIDE) Advanced Concept Technology Demonstration (ACTD)	This 5-year technology demonstration program focuses on integrating data from network management and information protection systems in order to provide automated, integrated, tactical warning and attack assessment. The program has set three objectives: • Develop an architecture for integration, analysis, and warning of IW attacks • Incorporate current and maturing intrusion sensing tools using expert systems technology for management of distributed systems • Correlate intrusion events at local, regional, and global command levels to improve the probability of detection and identification of IW events	The current implementation of the objective tool is EPIC^2. See the description of EPIC2 for information about the current properties of the objective tool.
Autonomous Agents for Intrusion Detection (AAFID)	This project is experimenting with a distributed architecture, within which various types of autonomous agents can be accommodated.	Contributing to defining the state of the art in distributed ADR systems.
Common Intrusion Detection Framework (CIDF)	This is a standards effort, by consortium, to develop protocols and application programming interfaces so that intrusion detection products can interoperate and components of them can be reused in other systems.	Defining the state of the art in this area.

Project	Thrust	Comment
DARPA Intrusion Detection Evaluation	This effort is to develop testing and evaluation standards; it is collecting and distributing the first standard corpus for evaluation of intrusion detection systems, both host-based and network-based.	Defining the state of the art in this area.
Distributed Intrusion Detection System (DIDS)	This effort was to develop intrusion detection for large, distributed networks employing multiple agents, both host-based and network-based, reporting to a centralized director, which would then be able to correlate inputs so that a single user could be tracked across multiple networks.	The information available on this project is from 1991. At that time, this effort was defining the state of the art. It has apparently since been overtaken by events.
Event Monitoring Enabling Responses to Anomalous Live Disturbances (EMERALD)	This project is attempting to develop a distributed scalable tool suite for tracking malicious activity through and across large networks. Using distributed monitors, it would provide a global detection and response capability to counter attacks occurring across an entire network enterprise.	The scale being attempted by this effort should raise issues that will help to refine the state of the art in systems that can span global networks. This effort is similar to others in progress and is a follow-on to NIDES.
Extensible Prototype for Information Command and Control ($EPIC^2$)	This project's goals are to achieve interoperability, integration, and coordination of intrusion control tools in a manner that can scale to global networks. The main thrust is to employ an expert system that can interact with diverse agents. Scalability is achieved by hierarchical arrangement of the expert systems.	Defining the state of the art in the use of expert systems.

Project	Thrust	Comment
Graph-based Intrusion Detection System (GrIDS)	This project is oriented toward detecting large-scale automated attacks on networked systems. The main idea is to build activity graphs in which nodes represent hosts and edges represent network activity among them. The detection technique is to compare a graph to a known pattern of intrusive activity.	Attempting to create a new detection method.
Next-Generation Intrusion Detection Expert System (NIDES)	This project created a system monitor, using a Sensors-Director architecture, employing both pattern matching and statistical deviation detection methods.	This system has been overtaken by events; see successor EMERALD.
Spitfire	This effort produced an intrusion alert manager implemented in a client/server architecture. The client: is a GUI providing access to data stored on the server. The server: provides access to an Oracle database that stores intrusion alerts, which can come from diverse sensors, and vulnerability and tool information.	This effort helped to push the state of the art in providing a user-friendly interface for managing intrusion alerts and in integrating vulnerability and tool (e.g., vulnerability scanners) information into the alert-operators environment.

Summary

We capsulize here the state of the art as we have reviewed it in this paper.

Table 6. Condensation of State of the Art in ADR

Topic	State
Detection reliability	Intrusion detectors can detect significant number of intrusions and short-term changes in behavior; but false-positives and false-negatives can still be a problem. The situation has improved, however, for run-of-the-mill attacks as the tools increasingly provide comprehensive coverage of known techniques and increasingly more frequent and user-friendly pattern updates. Vulnerability scanners can detect significant numbers of vulnerabilities; automatic updating of vulnerability libraries helps to alleviate the problem of detecting new vulnerabilities.
Detection of new attacks	Most IDSs detect only what they have been preprogrammed to detect; some progress in this area is now being made in COTS tools, some of which can detect variations on preprogrammed patterns; automatic updates to pattern libraries, which some vendors have begun to provide, also help in this area.
Reaction capability, including damage assessment and recovery, and real-time reaction	Current products are limited to providing input to the decision maker; there is no general capability for automated damage assessment and repair, but two commercial products do have some ability to fix illicit changes automatically (one can repair configuration changes, another uses a general response module that can interface to other management tools); useful automated response capability beyond notification depends on addressing several issues, for example how to ensure that the response system does not itself induce errors and how to safeguard against false positives and enemy-induced erroneous responses. Commercial tools, however, generally provide user interfaces that allow some customization, which may make it easier for an administrator or operator to better control these factors. Real-time reaction capability in monitoring tools is generally provided in the form of online alerting, e-mail notification, connection termination, and SNMP traps; a few tools can interact with firewalls; using SNMP traps, it should be possible to implement enterprise-specific reaction policies. Although responses in real time are limited in nature, they can occur rapidly enough to avoid or limit damage.

Topic	State
Analysis (e.g. consolidation and correlation of collected data)	COTS products have traditionally had little or no capability. This area has now been addressed by at least two vendors. For example, one vendor claims its product relates vulnerability data about hosts with attack data to show which hosts are both vulnerable and being attacked. In addition, a number of tools are able to collect data from distributed sources for presentation to the user. Several of the research projects we examined must employ some form of correlation technique to achieve their goals; the AIDE ACTD and EMERALD projects directly address attack-data correlation.
Traceback capability	This is an unsolved problem. Current tools may have traceback features, but they can be spoofed; traceback can require the cooperation of independent network-service providers, a sticky point.
Standards and terminology	Standards are lacking, but the CIDF and IETF[6] initiatives are promising. Terminology appears to be stabilizing with generally common understanding of terms.
Network topologies	One can generally expect that available tools can operate in Ethernet, Fast Ethernet, and FDDI environments. No tools deal directly with ATM.
Testing and validation	Testing and validation are inadequate at present but there is at least one initiative (DARPA Intrusion Detection Evaluation) to develop benchmarks for testing intrusion detection products and at least one vendor is providing specialized testing capability with its vulnerability scanning product.
Scalability	Scalability is lacking in some products for which scalability could be a feature (for many products, scalability is simply not an issue); some vendors (at least three) are beginning to address this area using distributed, hierarchical architectures; note also that the scale of applicability of tools in general can be quite large; some emphasis on scalability is apparent in the research projects we examined (EMERALD, EPIC2, and GrIDS).

6 Effective November 23, 1998, a new working group, called the Intrusion Detection Working Group, was formed in the Security Area of the IETF. The purpose of the Intrusion Detection Working Group is to define data formats and exchange procedures for sharing information of interest to intrusion detection and response systems, and to management systems which may need to interact with them.

Topic	State
Integration of tools into comprehensive ADR systems	Rudimentary capabilities are beginning to show up in COTS products, typically as a bundled suite of previously independent tools by the same vendor; this is an area that GOTS products developed earlier than the commercial vendors and current government-funded efforts continue to lead in this area.
Performance	False alarm rates can be high, especially if an intrusion tool is not tailored to the environment in which it operates, creating a performance problem for the analyst; many packet monitoring systems are not designed for networks operating at 100 megabits per second and higher.
Research and development in U.S. Government funded efforts	Correlation of data from several sources (e.g., EMERALD and GrIDS) and scalability (EMERALD, EPIC2, and GrIDS) are being worked on. There appears to be no research on damage assessment and recovery.
GOTS systems	These systems continue to be an important part of the protection arsenal for the military, especially for integrating tools and to achieve adequate scope, that is, the ability to deal with a large protection domain. This was apparently initially the case because commercial tools could not provide the capability needed; now, it appears to be more a matter of economics than technology.
Updates	Most vendors now provide updates for their intrusion detection and vulnerability scanning tools. In many cases, updates are product updates available to existing customers. In some cases, updates are issued on a regular basis, typically available via download or supplied on a floppy. In one case, automatic, periodic updates are available via the Internet.

List of References

1. Sandia National Laboratory, October 1996, *National Info-Sec Technical Baseline: Intrusion Detection and Response*, at URL: http://all.net/journal/ntb/ids.html on August 28, 1998, Lawrence Livermore National Laboratory, Sandia National Laboratory.

2. Hill, W. H., and S. J. Aguirre, September 1997, *Intrusion Detection Fly-Off: Implications for the United States Navy*, MITRE Technical Report 97W0000096, The MITRE Corporation, McLean, Virginia.

3. National Security Telecommunications Advisory Committee, December 1997, *Report on the NS/EP Implications of Intrusion Detection Technology Research and Development*, Network Group Intrusion Detection Subgroup Report, President's National Security Telecommunications Advisory Committee.

4. LaPadula, L. J., March 1999, *Compendium of Anomaly Detection and Reaction Tools and Projects*, Working Paper, Version 1.2, The MITRE Corporation, Bedford, Massachusetts.

5. Escamilla, T., 1998, *Intrusion Detection: Network Security Beyond the Firewall*, Wiley Computer Publishing, John Wiley & Sons, Inc., New York.

6. Hurwitz Group, Inc., March 1998, *Information Security: Assessing Risks and Detecting Intrusions*, White Paper viewable at Summit OnLine on November 2, 1998.

7. Amoroso, E. G., 1999, *Intrusion Detection: An Introduction to Internet Surveillance, Correlation, Traps, Trace Back, and Response*, Intrusion.Net Books, Sparta, New Jersey, http://www.intrusion.net/.

Appendix

About Anomaly Detection and Reaction Systems

The technology of anomaly detection and reaction is based on observation, experience, and classification of attacks, vulnerabilities, and countermeasures. Two key aspects of the technology today are the methods of detection and the types of tools available.

Methods of Detection

There are many terms used for describing methods of detection, but all the methods we have seen described fall into just one of two categories[7]

- Statistical deviation detection
- Pattern matching detection

Statistical Deviation Detection

In this approach the ADR tool looks for deviations from statistical measures. A baseline of values is defined for subjects and objects such as users, groups, workstations, servers, files, and network adapters. One can use historical data, simple counting, or expected values to establish the baseline. As activities being monitored occur, the ADR tool updates a list of statistical variables for each subject or object of interest. For example, the engine might count the number of files read by a particular user over a given period. This method treats any unacceptable deviation from expected values as an anomaly. For example, when the number of files read by a particular user over a given period exceeds the expected value for that period, the ADR tool declares a potential anomaly.

Practitioners use various terms for and explanations of this type of detection. Some are

- Anomaly detection: detecting deviation from a normal pattern of usage, as with an insider's use of an enterprise network

---

[7]In his recent book, Escamilla also identifies two types detection; he calls them statistical anomaly detection and pattern matching detection, referring to them as "IDS engine categories." We use the terms to encompass vulnerability scanners and policy compliance monitors as well. In his recent book, Amoroso discusses five methods used in practical intrusion detection; the five methods are not mutually exclusive and are often used in combination; this is a different and very useful way to view detection methods and Amoroso's discussion sheds light on how people and their automated tools actually operate [7].

- Statistical anomaly: detection is based on the assumption that users and networks exhibit predictable patterns of behavior from which they do not deviate significantly over short periods of time; a deviation from normal indicates a possible attack.
- Rule-based detection: detection based on a library of statistical descriptions of acceptable behaviors.

Pattern Matching

In this approach the ADR tool compares activity to stored patterns that model attacks or unacceptable states. Known attacks or types of attacks as well as proper configurations or system security policies, are modeled as patterns of data. Patterns can be composed of single events, sequences of events, thresholds of events, or expressions using AND, OR, and NOT operators[8]. This method treats any activity or state that matches a pattern as a potential anomaly.

Practitioners use various terms for and explanations of this type of detection. Some are

- Misuse detection: detecting attempted exploitation of a specific vulnerability
- Signature detection: detecting specific characteristics of a transmission or of the message being received
- Rule-based detection: detection based on a library of known attack patterns, unauthorized activity, or unacceptable system parameter values.
- Privilege anomaly[9] detection is based on analysis of audit logs. Analysis deduces an intrusion as follows: if a process has privilege "x" but there is nothing in the audit log showing how it got the privilege in an authorized manner, the analyzer deduces that the privilege was attained in an unauthorized manner.

Use of the Methods

The two methods of anomaly detection can be applied to both network-based detectors and host-based detectors. Network-based detectors normally are sited at a place in the network at which all the traffic going into and coming out of a domain of protection can be

---

8 As Escamilla notes, negation could be used for detecting unacceptable events but it might introduce computational complexity since it could require looking for "everything but this event" [5]. More likely, negation will be useful for modeling data values that are out of bounds.

9 This term is an exception to our claim that we are using the dialect of intrusion detection science; we have had to invent this term for the type of detection reported here, which is based on work done at Lincoln Lab.

monitored. Network traffic being the only input to a network-based detector, the detector does not depend on data from or cooperation with other systems on the network. A host-based detector, normally sited on the host it is to protect, can look for attacks related to specific vulnerabilities of its host, can analyze local system logs, and can monitor local user activity. However, host-based tools may use the Agents-Director architecture, in which case the Agent performs the collection function (e.g., gather data from event logs), may do some preliminary analysis or filtering of the data, and transmits the selected data to the Director, where analysis, correlation, storage, and other functions may be performed.

Types of Tools

Two traditional terms for intrusion detection systems are "host-based" and "network-based." These terms are likely to disappear from usage as the variety of anomaly detection and reaction systems available today no longer fit this simplistic taxonomy. More current terminology has been used by Escamilla, who identifies three main categories of intrusion detection tools [5].

- Scanners: A scanner is an IDS that periodically looks for vulnerabilities that might open a system to exploitation by an intruder or for evidence of intrusions after they have occurred. Scanners can run directly on the target or can scan targets from a remote location.
- System-level monitors: System-level monitors look for evidence of misuse and intrusion in real time. They may use pattern matching or statistical deviation detection. System-level monitors gather data from the target systems and typically send the data to a central director component that analyzes the data.
- Network sniffers: A network sniffer examines all incoming network packets on its subnet looking for indications of intrusion attempts. Typically, a network sniffer runs on a computer on the subnet whose network adapter operates in promiscuous mode[10]. A network sniffer can also run on some remote computer and communicate with a router that provides it network traffic information for the subnet being sniffed[11].

We can distinguish between scanners that look for vulnerabilities and scanners that look for infractions. Some practitioners mean both vulnerabilities and infractions when they use the term "vulnerabilities." That is, an infraction, such as a violation of access policy by an intruder

---

[10] In promiscuous mode, a network adapter sees all packets on its subnet. In normal usage, such an adapter captures and passes up to its host those packets intended for its host.

[11] NetRanger, by Cisco, provides a configuration such as this, in addition to being able to run in standalone mode with a promiscuous mode adapter.

who gains root access to a UNIX system, is looked upon as a vulnerability. Here the vulnerability is represented by the changes in the system caused by the infraction. The assumption is being made that the infraction actually did cause one or more changes in the system. However, the assumption may not always hold. The infraction of the example just given may compromise data through a read operation, leaving everything as it was before the infraction, except for entries that may have been made in logs and updates to system data such as "last time accessed" data for the file that was read. The change in "last time accessed" data most likely does not represent a vulnerability in the system. We will distinguish between scanning for vulnerabilities and scanning for infractions. An infraction scan could discover that an intrusion occurred by examining audit logs, even though there may be no other direct evidence of the intrusion.

A difference between an infraction scanner and a system-level monitor, as pointed out by Escamilla, is that the scanner operates periodically while the monitor operates in real-time. Also, one may detect intrusions not noticed by the other.

Although in theory a network sniffer could operate in nonreal time, in practice, sniffers operate in real time. We will treat all sniffers as operating in real time unless we come across a product that does not.

We use an extended terminology to characterize the type of an ADR product based on the foregoing and the following considerations:

- Some tools are integrated ADR systems. EPIC2 and ISS's SAFEsuite Decisions are examples. We will call such tools anomaly detection and reaction directors.
- We wish to include security compliance tools.
- We wish to distinguish analysis engines as an important class of tools.
- We wish to include support tools that provide data for anomaly analysis.

Thus, we recognize the following types of tools (listed alphabetically):

- Analysis Engine: An analysis engine receives inputs from a variety of sources (e.g., intrusion detectors, vulnerability scanners, etc.), possibly from widely distributed sources, and performs an analysis on the aggregated data to discover one or more things such as widely distributed attacks, distributed but coordinated attacks, patterns of vulnerabilities, etc.

- Anomaly detection and reaction director (ADRD or ADR Director): An ADRD integrates the functionality of two or more ADR tools; these tools may be of the same type or of different types. For example, an ADRD may integrate the functionality of many, identical network monitors or it may integrate the functionality of a system monitor and a vulnerability scanner. The ADRD provides an interface for managing ADR tools and their interactions. Products in this category may range widely in degree of integration. At a minimum, a system in this category provides a single interface to two or more instances of the same type of tool or to two or more types of tools that are interrelated at least via the view presented to the user. Very capable ADRDs include intrasystem communications among multiple instances and types of tools.
- Anomaly detection support tool: This kind of tool does not itself perform anomaly detection functions but gathers information that could be used to detect anomalies. Tools of this type might collect audit data from hosts or data from network packets, store the data in a database, and make it available in some user-friendly form.
- Infraction scanner: An infraction scanner periodically looks for evidence of infractions, including intrusions by outsiders and violations of policy by insiders.
- Network monitor: A network monitor looks for evidence of attempted misuse or intrusion in real time by examining data from network packets.
- Security compliance scanner: A security compliance scanner periodically examines the settings of system parameters that are relevant to the security of the system to ensure that they comply with a preset policy.
- System monitor: A system monitor looks for evidence of misuse and intrusion in real time by examining data from the target system.
- Vulnerability scanner: A vulnerability scanner periodically looks for vulnerabilities that might make a system susceptible to exploitation.

Architectural Attribute of Tools

One of the attributes used in the ADR Compendium for describing tools is called Architecture [4]. The ADR Compendium explains the possible values of this attribute as follows:

- Sensor: A Sensor is a software/hardware component that one adds to a system such as a server or workstation to provide anomaly detection and reaction functions specific to that system. An ADR Sensor can operate independently of other ADR capabilities to protect the system on which it is installed. It may also provide exported data or reports that can be used by other IDR capabilities. And, it may operate under the management of an ADR Director.

- Agent: An Agent is a software/hardware component that one adds to a system such as a router to provide anomaly detection and reaction functions specific to the domain of the Director under whose management it operates. An ADR Agent never operates independently; it is designed to work cooperatively with an ADR Director.
- Director: A Director is a software application or a software and hardware ensemble that performs storage, analysis, reporting, and/or command and control functions. It can be implemented on a stand-alone system or it can share a platform with other applications**,** running "independently" of the system on which it is installed, such as a server that hosts several different functions. An ADR Director controls or interacts with ADR Agents or Sensors within its domain. *See* description of ADR Director under Type of Tool above.
- Sensors-Director: self-explanatory
- Agents-Director: self-explanatory

Distribution List

Internal

B010

E. A. Palo
R. A. Games

D200

H. Gong

D320

R. W. Davis

D440

M. L. Kuras

D460

R. L. Dumas

D530

R. L. Pancotti

D550

B. W. Johnson

D620

B. D. Metcalf

G010

E. L. Lafferty
R. F. Nesbit

G020

D. J. Bodeau
W. R. Gerhart
H. W. Neugent
P. S. Tasker
J. M. Vasak

G021

C. L. Boeckman
L. J. Gill (Dept. File)
R. P. Galloni
J. D. Guttman
A. J. Jackman
L. J. LaPadula (5)
M. C. Michaud
P. D. Miller
S. L. Miravalle
J. Picciotto
M. L. Sheppard
R. J. Watro (5)

G022

J. G. DiChiara
J. L. Robinson (Dept. File)
R. W. Schmeichel, Jr.
J. R. Sebring
R. H. Walzer
J. T. Wittbold
M. M. Zuk

G023

S. J. Aguirre
M. S. Collins
G. J. Gagnon
F. C. Geck
S. M. Godin
T. A. Gregg
W. H. Hill, III
J. H. Tran
H. A. Robbins (Dept. File)

G025

K. M. Bitting
C. H. Bonneau
E. K Hawthorne
S. J. Moore
P. K. Townes (Dept. File)

G030

C. G. Greenbaum
J. P. L. Woodward
P. A. Pelletier

G037

R. Blount
D. Eyestone
T. R. Metcalf
T. Woodhouse

G038

S. P. Morrissey
L. S. Thomas

G045

J. F. Polito

G052

J. T. Rausch

W032

P. W. Attas
R. A. Duncan

Project

ESC/DIWS (co-located with ESC/DIG)
Hanscom AFB, MA 01731
Capt K. Fritz

ESC/DIW
Kelly AFB, TX 78226
R. Carter

ESC/DIWK
Kelly AFB, TX 78226
C. Durham

ESC/DIWP
Kelly AFB, TX 78226
R. Linaro

ESC/DIWS
Kelly AFB, TX 78226
C. Vera

ESC/DI
Hanscom AFB, MA 01731
M. Mlvezia

ESC/DIT
Hanscom AFB, MA 01731
D. Prince

Rome Operating Location of the Air Force Research Lab
AFRL/IFGB
525 Brooks Road
Rome, NY 13441-4505
D. Allain
J. Giordano
C. Maciag

External

AFIWC/EAC
250 Hall Boulevard, Suite 139
San Antonio, TX 78243-3143
M. Namatka
F. Ramirez

AFIWC/EAS
250 Hall Boulevard
San Antonio, TX 78243-3143
J. Cano

AFRL/IFGB
250 Hall Boulevard, Suite 214
San Antonio, TX 78243-3143
J. Pirog

Please do not delete these paragraphs or the final end-of-section mark in your document. They are important for correct functioning of the technical document template.
RoboTech: Version 1.0c