

AFRL-IF-RS-TR-2002-315
Final Technical Report
January 2003



**THE EMERALD MISSION-BASED CORRELATION
SYSTEM – AN EXPERIMENTAL DATA ANALYSIS
OF AIR FORCE RESEARCH LABORATORY
(AFRL) AIR FORCE ENTERPRISE DEFENSE
(AFED) INFORMATION SECURITY (INFOSEC)
ALARMS**

SRI International

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. J306


APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.


The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

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-2002-315 has been reviewed and is approved for publication.

APPROVED: 
GLEN E. BAHR
Project Engineer

FOR THE DIRECTOR: 
WARREN H. DEBANY, Technical Advisor
Information Grid Division
Information Directorate

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

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**
JANUARY 2003 **3. REPORT TYPE AND DATES COVERED**
Final Jan 02 – Jul 02

4. TITLE AND SUBTITLE
THE EMERALD MISSION-BASED CORRELATION SYSTEM – AN EXPERIMENTAL DATA ANALYSIS OF AIR FORCE RESEARCH LABORATORY (AFRL) AIR FORCE ENTERPRISE DEFENSE (AFED) INFORMATION SECURITY (INFOSEC) ALARMS

5. FUNDING NUMBERS
C - F30602-02-C-0024
PE - 62301E
PR - J306
TA - 01
WU - 08

6. AUTHOR(S)
Phillip Porras, Martin Fong, and Steven Chung

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
SRI International
333 Ravenswood Ave
Menlo Park California 94022

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Defense Advanced Research Projects Agency AFRL/IFGB
3701 North Fairfax Drive 525 Brooks Road
Arlington Virginia 22203-1714 Rome New York 13441-4505

10. SPONSORING / MONITORING AGENCY REPORT NUMBER
AFRL-IF-RS-TR-2002-315

11. SUPPLEMENTARY NOTES
AFRL Project Engineer: Glen E. Bahr/IFGB/(315) 330-3515/ Glen.Bahr@rl.af.mil

12a. DISTRIBUTION / AVAILABILITY STATEMENT
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 Words)
This project was established to experiment on the efficacy of the SRI EMERALD Mission-based Correlation System (M-Correlator) in analyzing INFOSEC device aloft in the Air Force Research Laboratory Information Directorate (AFRL/IF) Air Force Enterprise Defense (AFED) System. A large set of ISS RealSecure alerts produced within the AFRL network computing environment was analyzed by SRI using M-Correlator. Review of the M-Correlator experimental results identified a significant incident reduction capability, coupled with an effective alert ranking system. M-Correlator provided two orders of magnitude reduction in aloft, and affectively isolated highest-threat security incidents in the experimental date set. Further development may integrate a future M-Correlator release into the AFRL AFED system.

14. SUBJECT TERMS
Mission-Based Correlation, EMERALD, AFED, Defensive Information Warfare, Intrusion Detection, Alert Reduction

15. NUMBER OF PAGES
45

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

TABLE OF CONTENTS

Executive Summary	1
Goals and Expectations	2
M-Correlator Overview	4
A Framework for Better Alert Management	5
Mission-based Correlation	7
Pre-correlation Alert Processing	8
An Incident Handling Fact-Base	9
Incident Rank Calculation	11
Experiment Setup	15
Sensor Selection	15
Data Selection	16
AFRL AFED Mission Specification Definition	18
Topology Specification	18
Asset Criticality	19
Sensor Policy	20
Alert Aggregation Rules	22
M-Correlator Results	25
Results Summary	26
Example – The .69 Network SECADMIN VIEW	28
Example – The .69 Network OPERATOR VIEW	31
Observations and Lessons Learned	32
Uninteresting Subnetworks	33
The Power of Basic Aggregation	33
Incident Ranking Survives Sparse Alert Content	36
Future Work	38
Summary	39

FIGURES

Figure 1— M-Correlator Overview	6
Figure 2—Incident Rank Calculation	13
Figure 3—Topology Configuration	19
Figure 4—Critical Assets	20
Figure 5—SecAdmin View	21
Figure 6—Operator View	22
Figure 7—“Same IDS-Same Source-Same Class” Rule	23
Figure 8—.69 Network High Threat View	29
Figure 9—.69 Network Very Low Threat View	30
Figure 10—.69 Network Operator View	32
Figure 11—.42 Network Alert Type and Target IP Distribution	34
Figure 12 – .69 Network Alert Type and Target IP Distribution	36

TABLES

Table 1—Incident Handling Fact-Base Field Definitions	10
Table 2—Top Entries of 1,310 Subnets	17
Table 3—Summary of Experimental Data Analyses	27

EXECUTIVE SUMMARY

This Final Technical Report summarizes the activities performed during DARPA project F30602-02-C-0024, Mission-Based Correlation Experiment with AFRL AFED. This project began on January 18, 2002 and was completed on July 17, 2002. The project was established to experiment on the efficacy of the SRI EMERALD Mission-based Correlation System (M-Correlator) in analyzing INFOSEC device alerts in the Air Force Research Laboratory Information Directorate (AFRL/IF) Air Force Enterprise Defense (AFED) System. During this experiment, AFRL provided SRI a large set of ISS RealSecure alerts produced within the AFRL network computing environment. SRI performed an analysis of this data using M-Correlator and presented the results to DARPA and AFRL AFED personnel.

The review of the M-Correlator experimental results by AFRL AFED personnel identified a significant incident reduction capability, coupled with an effective alert ranking system. M-Correlator provided two orders of magnitude reduction in alerts, and effectively isolated highest-threat security incidents in the experimental data set provided by AFRL. Based on these results, the AFRL AFED group has requested that SRI produce a statement of work to extend the initial contract to provide an M-Correlator release that is integrated into the AFRL AFED system.

GOALS AND EXPECTATIONS

In April 1999, SRI's EMERALD development group began an effort under the DARPA-sponsored Cyber Panel research program to develop a mission-impact-based approach to the analysis of security alerts produced by distributed heterogeneous information security devices. This research led to the development of a prototype system called M-Correlator, which is capable of receiving security alert reports from a variety of INFOSEC devices. Once translated to an internal incident report format, INFOSEC alerts are augmented and, where possible, fused together under a correlated security incident report through a chain of processing. A *relevance score* is produced through a comparison of the incident target's known topology against the vulnerability requirements of the incident type, which is provided to M-Correlator by an *Incident Handling Fact-Base*. Next, a *priority calculation* is performed per incident to indicate (a) the degree to which the incident has targeted at critical assets and (b) the amount of interest the user has registered for this incident type. Finally, an overall *incident rank* is assigned to each incident, which brings together the priority of the incident with the likelihood of success.

Under the DARPA Cyber Panel program, M-Correlator reached a level of maturity capable of supporting experimental trials on live operational data sets. In light of this progress, DARPA selected M-Correlator for testing in a state-of-the-art INFOSEC alert management environment, the Air Force Research Laboratory (AFRL) Air Force Enterprise Defense (AFED) System. The objective of this project is to demonstrate the efficacy of the M-Correlator prototype as a security-incident correlation service for use by AFED analysts. SRI applied its M-Correlation algorithm to eight weeks of operational AFED alert data, and demonstrated a significant incident reduction capability coupled with an incident ranking system that automatically isolated the highest-threat security incidents within the M-Correlated incident database. In this experiment, M-Correlator

demonstrated a strong algorithmic basis for mission representation in the context of INFOSEC alert management, and employed this mission representation as a basis for quickly removing from consideration low-interest, false positive alerts, without the use of filters that may lead to data loss.

Another key facet of this experiment was to understand the practical issues of M-Correlator deployment, by exploring the issues of building mission specifications and topology databases in support of large datasets from real operational networks. The results from this experiment were presented in detail to AFRL AFED personnel, to gain an end user's perspective of the value of M-Correlator results. AFRL AFED personnel were impressed with the results to a degree that they are seeking to establish follow-on funding that will allow M-Correlator to be permanently deployed within the AFRL AFED system.

M-CORRELATOR OVERVIEW

Among the most visible areas of active research in the intrusion detection community is the development of technologies to manage and make sense of the growing availability of security-relevant alert streams that are produced from the increasing deployment of INFOSEC devices. The motivation for Mission-based Correlation is straightforward: as we continue to incorporate and distribute advanced security services into our networks, we need the ability to understand the various forms of hostile and fault-related activity that our security services observe as they help to preserve the operational requirements of our system. Today, in the absence of significant fieldable technology for security-incident correlation, several challenges face those attempting to provide effective security management in mission-critical network environments:

- Domain expertise in understanding and isolating the highest threat activity that is encountered daily by an active and visible Internet-connected network is not widely available. Also not widely available are skills in understanding under what conditions one can merge INFOSEC alerts from different sources, such as merging firewall and OS syslogs with intrusion detection reports. In an environment where hundreds (or even thousands) of INFOSEC alarms may be produced daily, it is difficult to understand redundancies in alert production that could simplify the interpretation of the alerts, and there is no technology to help prioritize which security incidents pose the greatest threat to one's administrative responsibilities.
- The sheer volume of INFOSEC device alerts makes security management a time-consuming and therefore expensive effort. For example, one financial institution that attempted to deploy ISS RealSecure found even a small deployment of RealSecure sensors to be an overwhelming management cost. As a result, these intrusion detection components were tuned down to an extremely narrow and ad hoc selection of a few detection heuristics, effectively minimizing the

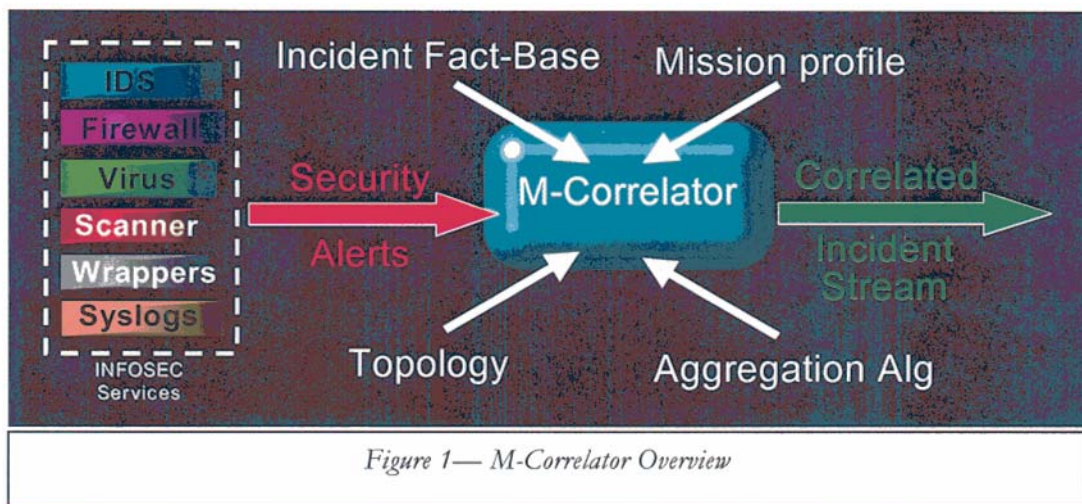
coverage of the IDS tool. In addition, the planned large-scale IDS deployment by the financial institution was postponed and ultimately scaled down.

- In managing INFOSEC devices, it is difficult to leverage potentially complementary information that would allow better interpretation of hostile activity or system distress, and could allow greater alert reduction through alert clustering or aggregation. As a result, security-relevant information that, for example, is captured in a firewall log is typically manually analyzed in isolation from potentially relevant alert information captured by an IDS, syslog, or other INFOSEC alert source.
- With respect to cross-domain, or enterprise security management, there are virtually no tools fielded to correlate security incidents reported in one domain with security incidents observed in other administrative domains. As a result, vital information regarding security attacks from hostile external agents may not be shared, limiting the ability for coordinated response or advanced warnings to the presence of wide-scale threats.

A FRAMEWORK FOR BETTER ALERT MANAGEMENT

The Mission-based Intrusion Report Correlation System, or *M-Correlator*, is designed to consolidate and rank a stream of security incidents relative to the needs of the M-Correlator operator, given the topology and operational objectives of the protected network. Figure 1 illustrates the conceptual elements of the M-Correlator system. The M-Correlator is capable of receiving security alert reports from a variety of INFOSEC devices, such as firewalls, intrusion detection systems, and antivirus software. The following briefly summarizes the M-Correlator processing algorithm. In subsequent sections, each step of processing is described in detail.

Once translated to an internal incident report format, INFOSEC alerts are augmented and, where possible, fused together through a chain of processing. The first phase of processing involves dynamically controllable filters, which can provide remote subscribers to INFOSEC alerts with the ability to eliminate low-interest alerts, while not preventing the INFOSEC devices from producing those alerts that may otherwise be of interest to other administrators. Next, the alerts are vetted against the known topology of the target network. A *relevance score* is produced through a comparison of the alert target's known topology against the known vulnerability requirements of the incident type, which is provided to M-Correlator by an *Incident Handling Fact-Base*, which represents a variety



of critical information per alert type. Next, a *priority calculation* is performed per alert to indicate a) the degree to which the alert is targeted at critical assets and b) the amount of interest the user has registered for this alert type. Lastly, an overall *incident rank* is assigned to each alert, which indicates

- The degree to which the incident appears to impact the overall mission of the network, as reflected in the priority calculation
- The probability that the activity reported in this alert was successful, as derived from the relevance score and, when available, alert outcome attributes provided by the INFOSEC device

Once ranked, the M-Correlator attempts to combine related incident alarms with an attribute-base *alert clustering algorithm*. The resulting correlated incident stream represents a filtered, lower-volume, content-rich security-incident stream, with an incident ranking scheme that allows M-Correlator operators to identify those incidents that pose the greatest risk to the currently specified mission objectives of the monitored network.

MISSION-BASED CORRELATION

The objective of M-Correlation is to fuse related alerts into higher-level security incidents, and rank them based on the degree of threat each incident poses to the mission objectives of the target network, as specified by the correlator operator. To understand the process of security-incident ranking, we must first define the notion of *mission* with respect to an administrative network domain. An underlying assumption of this methodology is that every network-computing environment has been assembled with some degree of expectation that it provides important services and/or data access to the users of that network. The degree to which hosts and data resources contribute to the objectives of the user community will vary.

The *mission* is the underlying objective for which the computing resources and data assets of the monitored network are brought together and used. Operators provide the M-Correlator a *mission specification* of the protected network, through which they register

- Critical computing assets (such as file servers on which the user community depends)
- Critical network services (such as Web server, DBMS)
- Sensitive data assets (these are primarily files and directories considered highly sensitive or important to the mission of the network)

- Administrative accounts and nontrusted user accounts such as might be used by consultants

In addition, the expression of *mission specification* also includes interest indicators provided by the M-Correlator operator based on incident type or class. The M-Correlator defines ten incident classes derived from an analysis of previous work in alert classifications and taxonomies, as well as from practical experience gained by SRI in analyzing INFOSEC alarms. For each incident class, the operator may specify high, medium, or low interest, or may indicate interest level by specifying a discrete weighting from 0 to a maximum of 255.

PRE-CORRELATION ALERT PROCESSING

SRI has been actively involved in advanced research technology for host and network intrusion detection, intrusion response, IDS interoperability, and intrusion report correlation. In the area of IDS interoperability, SRI has participated extensively in foundational work in the DARPA-sponsored Common Intrusion Detection Framework (CIDF), where SRI led the initial development of the CIDF Common Intrusion Specification Language (CISL). CISL is a general language, capable of expressing a wide range of observations, conclusions, and response prescriptions. Subsequently, SRI co-presented concepts and requirements at the initial meeting of the IETF working group on IDS interoperability, which has most recently led to the first Intrusion Detection Message Exchange Format (IDMEF) specification.

During the last six years, SRI has gained deep insight into the limitation of INFOSEC alert content, and experience in extracting INFOSEC alerts from security products such as firewalls and intrusion detection systems. SRI has had the unusual problem of having to provide a consistent record structure to service reports produced from a variety of algorithms and event streams. SRI's EMERALD sensors employ a rule-based signature analysis engine called eXpert, a statistical

anomaly detection engine called eStat, or a Bayesian network-based analysis engine called eBayes. EMERALD sensors have also been instantiated to analyze multiple event streams at the host audit, application log, and network traffic layers. Sensors must also operate in real time or batch mode. This wide variety of requirements has led SRI to develop a general-purpose unified alert and correlation-reporting format that is highly extensible to support the myriad of general-purpose INFOSEC devices and correlation services.

AN INCIDENT HANDLING FACT-BASE

The SRI M-Correlator prototype includes an *Incident Handling Fact-Base* that provides the M-Correlator the necessary input to recognize the configuration dependencies of security incidents, clustering information to help the M-Correlator better fuse alerts and recognize equivalent incidents, and incident classification information to help M-Correlator operators preselect alerts for filtering and register their interest levels in certain problem activity. The incident handling fact-base provides critical information needed to

- Augment terse INFOSEC device alerts with meaningful descriptive information, associate alerts with M-Correlator specific incident codes and classifications, and provide counteraction recommendation for use by visualization and response tools
- Understand the dependencies of incident types on OS type and version, hardware platform, available network services, and applications
- Understand which incident types can be merged by the M-Correlator alert clustering algorithm

Table 1 provides the field definitions of entries in the M-Correlator incident handling fact-base. Entries in this fact-base are referenced in subsequent sections, which describe topology vetting, prioritization, incident ranking, and alert clustering. The current M-Correlator fact-base provides

incident definitions for more than 1,000 known intrusion reports from ISS RealSecure, Snort, the EMERALD suite of host and network-based intrusion detection sensors, and Checkpoint’s Firewall-1 product line. Incident types that are not represented in this fact-base can still be managed and aggregated by the M-Correlator; however, the advanced alert clustering and relevance calculations are not performed on alerts that are absent from this fact-base.

Table 1—Incident Handling Fact-Base Field Definitions

FIELD TYPE	DESCRIPTION
Incident Code	A unique code to indicate incident type. These codes have been derived from the original Boeing/NAI IDIP incident codes that were used by the Common Intrusion Detection Framework CISL specification. A mapping between this incident code and other well-known attack code specifications such as Bugtraq ID, CERT ID, and MITRE CVE codes is available using the References field (below). A mapping to commercial-specific codes is available using the COTS Codes field.
COTS Codes	Equivalent codes of well-known commercial off-the-shelf (COTS) incident name or numeric code value that express this incident.
Incident Class	An M-Correlator general categorization scheme used for abstractly registering interest in an incident that represents a common impact on the system. Incident types are associated with only one incident class. The following incident classes are defined in the M-Correlator Fact-base: Probe, Access-Violation, Integrity-Violation, System-Environment-Corruption, User-Environment-Corruption, Asset-Distress, Suspicious-Usage, Connection-Violation, Binary-Subversion,

	Action-Logged.
Description	Human-readable incident description.
Directives	Used by response component and not referenced by M-Correlator. This field provides recommended response directives. There are currently eight response directives defined: Diagnose, Lockout, Kill, Checkcfg, Fixprems, Filter, Notify, and Reset.
Vulnerable OS and Hardware	OS types and version and hardware architectures that are required for the successful invocation of the incident.
Bound Ports and Applications	Required network services and applications for this incident type to succeed.
Cluster List	One or more index values that may be associated with incident types. Two alerts that share a common cluster name may be candidates for merger should other attributes be aligned.
References	Bugtraq ID, CERT ID, Common Vulnerabilities and Exposures (CVE) ID, available descriptive URL.

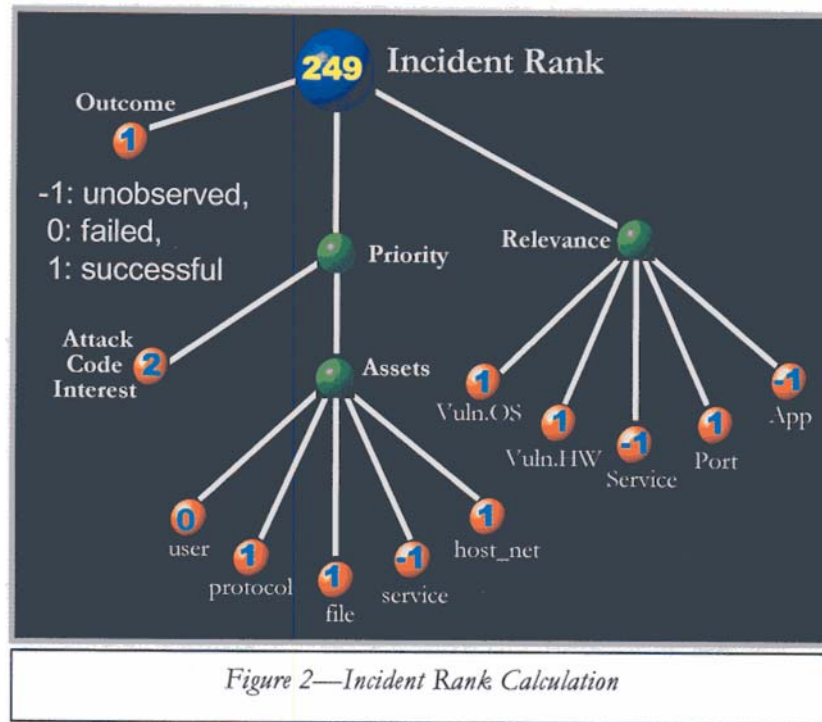
INCIDENT RANK CALCULATION

At the heart of the M-Correlator is a mission-based incident ranking algorithm, which ranks the incoming stream of diverse INFOSEC alerts relative to the needs of the M-Correlator operator, as

expressed in the *mission specification* and topology of the protected network. In environments where literally thousands of INFOSEC alerts are produced on a daily basis, automated incident ranking provides one critical strategy to help identify the highest-threat security incidents. Incident ranking is intended to augment the security expertise of network administrators who may not continually maintain in-depth knowledge of security alerts and their possible impact on the network mission. Equally important, incident ranking helps lower the cost of security administration by helping to recognize and label lower-threat incidents. This allows operators to pay less consideration to low-threat events, while still making these events available in the security database should the need arise, (i.e., incident ranking allows one to cost-effectively manage INFOSEC alerts without disabling or “tuning down” INFOSEC alert reporting services, which is common practice today). The M-Correlator breaks the incident rank calculation into three parts, as shown in Error! Reference source not found.:

- **Relevance Calculation:** M-Correlator maintains an internal topology map of the protected network, which is dynamically managed by the user (an *automated* topology analysis could be supported). M-Correlator’s topology map indicates the operating system and version, platform computing hardware, (e.g., Sparc, Intel,) available network services and port bindings, (e.g., FTP:21-tcp, SSH:22-tcp, HTTP:80-tcp,) and the list of mission-critical applications running on the host, (e.g., DBMS). As each alert is processed by the M-Correlator, the associated list of known dependencies for that alert, as indicated within the incident handling fact-base, is compared against the configuration of the target machine. Positive and negative matches against these required dependencies will result in increased or decreased weighting of the relevance score, respectively. An unknown alert, incompletely specified dependency information in the fact-base, or incomplete topology information regarding the target host, will result in a neutral

weighting of the relevance score (i.e., the score will not contribute positively or negatively to the overall incident rank).



- Alert Priority Calculation:** The alert priority calculation is formulated through a comparison of the alert attributes against the attributes of the mission specification. The alert priority calculation is formulated from two contributors: an *attack code interest* and the overall *criticality score* of the target asset. The attack code interest is an operator-provided weighting of the importance of each incident class (operators can also specify interest on a per-incident basis). Currently, the M-Correlator allows operators to specify low, medium, or high interest, or operators may weight their interest in an incident class on a scale from 0 to 255. The criticality score represents the degree to which a reported attack is found to target mission-critical assets, network services or applications, to affect critical user accounts, or to target mission-critical resources (e.g., critical files or directories).

- **Incident Rank Calculation:** Incident ranking represents an assessment of each security incident from 0...255 (0 representing a low score, 128 representing a neutral score, and 255 representing the highest incident rank). Incident rank formulation is performed with respect to a) the incident's impact on the mission profile as reflected by the M-alert priority calculation operator and b) the probability that the security incident reported by the INFOSEC device(s) has succeeded. Figure 2 illustrates how incident rank is formulated from the relevance calculation, mission priority, and an alert outcome attribute (where provided by some INFOSEC devices). Most IDS sensors provide little if any indication regarding the outcome of an observed intrusion event. Therefore, sensor outcome attributes are augmented by the relevance calculation, which indicates the M-Correlator's assessment as to whether the system was susceptible to the attack.

EXPERIMENT SETUP

The initial phase of the experiment involved the normalization of AFRL AFED sensor alert content and augmentation of the M-Correlator Incident Handling Fact-Base. During this experiment setup phase, SRI examined a large corpus of alerts produced within Rome Laboratory's operational computing network for use in the M-Correlator analysis. SRI selected one sensor's content from which to perform the M-Correlator experiment, and selected three network domains with rich alert diversity from which to explore mission-based impact analysis.

SENSOR SELECTION

The data set used in this experiment pertains to reports generated by a suite of intrusion detection systems that monitored a large enterprise network within the U.S. Air Force. The IDSs deployed were ISS RealSecure, NetRadar, and a CMU anomaly detection sensor. The output of these IDSs was put into an AFED database. The database records provide event insertion time, sensor name, source/destination IP address, attack signature, and source/destination port number. The data in the event table of the AFED database was used as the raw input for this experiment.

The configuration information of these sensors was not available to this experiment. The sensors were apparently deployed in the "peering point" mode, and the monitored traffic included both internal-internal traffic and internal-external traffic. From the data set, SRI identified 51,725 distinct destination IP addresses. Based on the first three octets of these addresses, SRI partitioned them into 1,310 groups. Each group corresponds to a logical class C subnet and includes all IP addresses that have the same three-octet prefix. Because SRI did not have access to the network topology information, these subnets may or may not correspond to the physical networks within AFED.

Moreover, the majority of these destination addresses may not correspond to physical machines; they were identified from the IDS alerts, and thus may merely correspond to targets of sweeps and probes.

The data set consists of 737,543 alerts. Among these alerts, 731,346 of them were generated by RealSecure, 5,476 from NetRadar, and 721 from the CMU sensor. The earliest timestamp of the alerts is in 7/23/2001 and the latest one is in 1/9/2002. Based on SRI's examination of the data set, individual sensors appeared to operate within a smaller time period than indicated by those timestamps. In particular, alerts from RealSecure were from 8/8/2001 to 10/5/2001, those from NetRadar were from 7/23/2001 to 1/9/2002, and those from the CMU sensor were from 9/28/2001 to 12/13/2001. Moreover, there were "holes" in detection coverage during which no alerts were produced from SRI, which is highly indicative of sensor downtime. During the periods when RealSecure was active, it averaged approximately 21,000 alerts per day.

From the incident signature point of view, the data set has 125 distinct signatures. Among these signatures, 97 were from RealSecure, 18 were from NetRadar, and 10 were from the CMU sensor. The distribution of these alerts is quite skewed. For instance, the most popular incident signature, namely, IRC, accounted for 600,172 alerts.

DATA SELECTION

Because of a lack of information for the semantics of the NetRadar and the CMU sensor signatures and for the network topology information, SRI decided to select a subset of the data set and used it in the experiment. The selection involves focusing on RealSecure alerts and on alerts pertaining to certain subnets.

SRI used the following criteria in selecting the subnets: number of alerts per subnet, number of hosts per subnet, and number of distinct alert types per subnet. An interesting input data set from the perspective of alert correlation has the following properties: it has a large number of alerts, the number of hosts involved is reasonably large, and the alerts have some diversity.

SRI compiled a table for the 1,310 subnets based on the above-mentioned criteria. Table 2 shows the top entries of the table in the decreasing order of the number of alerts per subnet. Based on the results, SRI selected three subnets, namely, 1xx.1xx.69, 1xx.1xx.129, and 1xx.1xx.42. For each of these subnets, SRI developed a hypothetical mission specification that includes network topology, host/service criticality, aggregation criteria, and incident interests. These specifications were used by M-Correlator in ranking and aggregating alerts.

Subnet	Num. Alerts	Num. Hosts	Distinct Signatures
1xx.2xx.192	124227	4	1
1xx.2xx.244	119003	1	1
1xx.1xx.069	25624	49	26
1xx.1xx.129	21896	47	13
1xx.2xx.002	12189	251	2
1xx.2xx.007	10185	222	2
1xx.2xx.017	9640	209	2
1xx.2xx.231	8615	2	1
1xx.2xx.301	7791	4	1
1xx.2xx.086	7463	4	1
1xx.2xx.082	6658	3	1
1xx.2xx.199	5954	5	2
1xx.2xx.059	5731	89	2
1xx.2xx.075	5583	245	2
1xx.2xx.183	5318	4	2
1xx.1xx.042	5310	17	30

Table 2—Top Entries of 1,310 Subnets

AFRL AFED MISSION SPECIFICATION DEFINITION

Once alert selection and content normalization was completed, the next phase of the experiment involved the development of a mission specification and topology database. Because of logistical issues and availability of AFRL AFED personnel, both sides agreed that for this experiment, SRI would develop hypothetical mission specifications that would, as realistically as possible, resemble mission specifications that could be deployed and used in live AFRL AFED operations. In addition, because of the sensitive nature of the internal Rome Laboratory computing network, it was decided that SRI would develop hypothetical topology database content in performing the M-Correlator relevance calculation. The following briefly summarizes the topology definition of the selected networks, the asset criticality enumerations, the sensor policy, and the alert aggregation algorithms used during the AFRL AFED experiment.

TOPOLOGY SPECIFICATION

Based on the IP address, port number, and attack signature information extracted from the alerts, SRI developed a hypothesized topology configuration for the hosts in that subnet. The topology information can usually be obtained by using some network management tools such as *nmap*. Because the topology information was not available to this project, SRI had to rely on its experiences in topology development, along with strong clues provided in the raw RealSecure alert stream, to build a realistic topology database.

For example, the following topology configuration for the 1xx.1xx.42 subnet specifies that the host with IP address 1xx.1xx.42.14 is a FreeBSD box and runs the specified TCP and UDP services (e.g., FTP and NTP). The topology information is used to calculate the relevance score of an alert. For

instance, alerts corresponding to attacks that are specific to Windows boxes against a FreeBSD box or to attacks against a service that the target host does not provide will receive a low relevance score.

```
Topology_config
{
    ...
    1xx.1xx.42.14
    {
        Hostname           WebServer14
        OperatingSystem    FreeBSD
        Hardware           Intel
        BoundTCPServiceList [ ftp ssh telnet smtp finger http
                               sunrpc login shell submission ]
        BoundUDPServiceList [ sunrpc ntp biff syslog ntalk
                               urm nqs sift-uft doom ]
    }
    ...
}
```

Figure 3 —Topology Configuration

ASSET CRITICALITY

A list of the hosts and services in the subnet was specified as critical in the mission specification. For example, the following critical asset specification was used for the 1xx.1xx.42 subnet. In this specification, the host 1xx.1xx.42.14 is specified as a critical host and the HTTP service it provides is also specified as mission critical. Alerts pertaining to critical hosts and critical services receive a higher priority score. In M-Correlator, one can also specify certain resources (e.g., files and

directories) as critical and certain user accounts as suspicious, and these attributes could also affect the priority score. Because the IDS reports used in this experiment did not provide that information, no critical resource and suspicious user information was specified in the hypothesized mission specification.

```
Critical_Assets  [  
    :1xx.1xx.42.14:80  
    :1xx.1xx.42.21:80  
    :1xx.1xx.42.59:80  
    :1xx.1xx.42.60:80  
    :1xx.1xx.42.14:53  
    :1xx.1xx.42.21:53  
    :1xx.1xx.42.59:53  
    :1xx.1xx.42.60:53  
    :1xx.1xx.42.14:  
    :1xx.1xx.42.21:  
    :1xx.1xx.42.59:  
    :1xx.1xx.42.60:  
    ]
```

Figure 4—Critical Assets

SENSOR POLICY

The priority score of an alert depends on the interest level assigned by the user to different classes of sensor alerts. A user that is more interested in availability than access violation may assign “High” interest level for the former and a “Low” interest level for the latter. This information is specified in the sensor policy section of the mission specification.

```

SecAdmin View:

    DEFAULT {
        PRIVILEGE_VIOLATION           High
        SYSTEM_ENV_CORRUPTION         High
        EXFILTRATION                   High
        ACCESS_VIOLATION               High

        INTEGRITY_VIOLATION           Medium
        BINARY_SUBVERSION              Medium
        USER_ENV_CORRUPTION            Medium
        USER_SUBVERSION                Medium

        PROBE                           MediumLow

        DENIAL_OF_SERVICE              Low
        ASSET_DISTRESS                 Low
        SUSPICIOUS_USAGE                Low
        CONNECTION_VIOLATION           Low
        ACTION_LOGGED                   Low

        RESERVED1                      Low
        RESERVED2                      Low
        RESERVED3                      Low
        RESERVED4                      Low
        RESERVED5                      Low
        DefaultClass                    Low
        MinAnomalyScore                 90
        MinConfidenceScore              09
    }

```

Figure 5—SecAdmin View

In the experiment, SRI developed hypothesized interest profiles for two types of users; one corresponds to a security administrator’s view and the other corresponds to an operator’s view. The SecAdmin view specifies a high interest in privilege violation, system-environment corruption, exfiltration, and access violation; a medium interest in integrity violation, binary subversion, user-environment corruption, and user subversion; a medium-to-low interest in probe; and low interest in denial of service, asset distress, suspicious usage, connection violation, and action logged.

The operator view specifies a high interest in denial of service, asset distress, system-environment corruption, integrity violation, and binary subversion; medium-to-low interest in privilege violation;

and low interest in the other incident classes. As illustrated later, using these two different user interest profiles gave rise to different alert ranking results in the experiment, reflecting the kind of activities considered to be more important to the user.

Operator View:

DEFAULT {	
DENIAL_OF_SERVICE	High
ASSET_DISTRESS	High
SYSTEM_ENV_CORRUPTION	High
INTEGRITY_VIOLATION	High
BINARY_SUBVERSION	High
PRIVILEGE_VIOLATION	MediumLow
ACCESS_VIOLATION	Low
EXFILTRATION	Low
USER_ENV_CORRUPTION	Low
USER_SUBVERSION	Low
SUSPICIOUS_USAGE	Low
PROBE	Low
CONNECTION_VIOLATION	Low
ACTION_LOGGED	Low
RESERVED1	Low
RESERVED2	Low
RESERVED3	Low
RESERVED4	Low
RESERVED5	Low
DefaultClass	Low
MinAnomalyScore	90
MinConfidenceScore	09
}	

Figure 6—Operator View

ALERT AGGREGATION RULES

M-Correlator is equipped with an alert aggregation language that allows users to specify alert fusion logic. In an aggregation rule, one specifies how long M-Correlator will keep a meta-alert (i.e., aggregated alert) when there are no subsequent alerts that can be aggregated with it (cf. the *Delay*

Until Expire parameter); the delay before a report for a meta-alert is initially issued (cf. the *Initial Flush Delay* parameter); and the refresh interval between successive meta-alert updates (cf. the *Delay Until Flush* parameter). When an alert can be aggregated with a meta-alert, the *Unique Match* attribute of an aggregation rule determines whether this incoming alert will be considered by other aggregation rules. Moreover, the *Merge Action* attribute determines how to combine the data in the incoming alert with that of the meta-alert. Finally, the *Match If* part specifies the conditions that need to be met by the incoming alert and the meta-alert to trigger the aggregation.

```

{
  Profile           Same_IDS_Same_Source_Same_Class
  Policy           Liberal
  Delay_Until_Expire 1800
  Delay_Until_Flush 90
  Initial_Flush_Delay 90

  Enable           true
  Unique_Match     false
  Merge_Action     fuse

  Match_If [ AND
    [ EQ           observer_ID
    [ OVERLAP      source_IParray ]
    [ EQ           incident_class ]
    [ OR
      [ EQ         incident_signature ]
      [ EQ_CONST  incident_class 4 ]
      [ OVERLAP    target_TCP_portarray ]
    ]
  ]
}

```

Figure 7 — “Same IDS-Same Source-Same Class” Rule

In this experiment, SRI deployed two aggregation rules, namely, *Same IDS Same Source Same Class* and *Incident Class and Dest Match*. The former is attack-source-oriented and specifies the following conditions for aggregation: the incoming alert and the meta-alert carry the same observer ID, (i.e., the activities are observed by the same sensor,) have overlapping source IP addresses, (i.e., aggregating activities from the same source,) and belong to the same incident class. To ensure that

the aggregation performed by this rule was not too aggressive, (e.g., aggregating multiple attack attempts from the same source against different network services on different hosts to gain remote access,) SRI added a condition to the rule to restrict its power. Specifically, the incoming alert and the meta-alert must satisfy one of the following criteria:

- (1) They have the same incident signature (thus they correspond to the same attack).
- (2) They are both probes (thus they both correspond to reconnaissance activities).
- (3) Their target TCP port arrays overlap (thus they are both after the same network service).

The second aggregation rule used in this experiment, *Incident Class And Dest Match*, is attack target oriented. Distributed attacks from different sources against the same target could be aggregated by this rule. Basically, it aggregates alerts that target the same host and belong to the same incident class. SRI also added a condition to this rule to restrict its aggressiveness by specifying that the incident class must be one of the following: denial of service (class 3), probe (class 4), or asset distress (class 9).

M-CORRELATOR RESULTS

SRI performed a series of experimental runs using the ISS RealSecure alert repository provided by Rome Laboratory. As discussed in Section *Data Selection*, SRI isolated the RealSecure alerts of three subnetworks, and developed two independent mission specifications (SecAdmin and Operator) as described in Section *Table 2—Top Entries of 1,310 Subnets*

AFRL AFED Mission Specification Definition. SRI performed six independent experimental data analysis runs. For each selected subnet, SRI ran the network's associated alerts through M-Correlator by using each of the two mission specifications.

The results of each data analysis run were written directly from M-Correlator to an RDBMS (postgres was used in this experiment). Once correlated reports are stored into an EMERALD database, the EMERALD Alert Management Interface (eAMI) version 2.0 is used to display and organize M-Correlator's ranked security incident reports into folders. eAMI provides an effective means of displaying the synthetic correlation attributes created by M-Correlator in the following discussion various screenshots from eAMI will be used to present M-Correlator results.

The analyses of the M-Correlator results are presented to emphasize three aspects of operation. Only two of these aspects are viewable directly through eAMI displays. For M-Correlator to demonstrate itself as a valuable and practical tool for deployment, the following issues were considered:

- (4) Ease of configuration — SRI demonstrated an ability to create, with minimal support from AFRL AFED personnel, two realistic mission specifications. However, this aspect of operational use was not fully explored, as AFRL AFED personnel did not independently

modify or create a mission specification. In the next phase of this experiment, SRI is planning to deliver a release of M-Correlator into the AFED system, and will train AFRL AFED personnel on configuration and mission specification development.

- (5) Effectiveness in alert reduction — SRI demonstrated an ability to provide two to three orders of magnitude reduction in the raw alert count. SRI further presented evidence to AFRL AFED personnel of content merging and showed how correlation provides minimal content loss from the raw ISS RealSecure alerts.
- (6) Effectiveness of incident ranking scheme — SRI demonstrated a highly effective ranking system that provided significant separation between high and low threat alerts. SRI and AFRL AFED personnel co-reviewed the High Threat folder content and confirmed that each alert in the High Threat folder was appropriately classified, given the mission specification definition. Both groups also examined the lower threat incidents and agreed that these alerts were appropriately classified.

RESULTS SUMMARY

Table 3 summarizes alert and incident counts produced by the M-Correlator Experimental data analyses. Columns 2 through 4 of Table 3 summarize the results produced for each of the selected subnets. Three networks were involved in the experiment, the 128.132.42.* subnet, the 128.132.69.* subnet, and the 128.132.129.* subnet.

Row 1 of Table 3 summarizes the total number of raw ISS RealSecure alerts analyzed by M-Correlator, in which the source or destination IP address was found to match the IP address space of the selected subnet. Row 2 indicates the number of unique ISS alert types that were reported for the selected subnet. Row 3 identifies the greatest variety of alert types at 30. Row 4 identifies the greatest number of alert types at 15. Row 5 identifies the greatest number of incidents at 10. Row 6 identifies the greatest number of incidents at 10. Row 7 identifies the greatest number of incidents at 10. Row 8 identifies the greatest number of incidents at 10. Row 9 identifies the greatest number of incidents at 10. Row 10 identifies the greatest number of incidents at 10.

the number of unique IP addresses in the selected subnet that were involved, as either the source or target, in a raw security alert provided by RealSecure.

	.42 Network	.69 Network	.129 Network
Total Number of Alerts	5,310	25,624	21,896
Total Unique Alert Types	30	26	13
Total Unique Target Hosts	17	49	47
Aggregated Security Incidents	357	250	78
Privilege alerts	291	8	3
Bad access alerts	6	2	1
Availability alerts	14	33	18
Suspicious alerts	26	182	42
Probe alerts	20	25	14
SECADMIN Incident Ranking			
High Threat	133	4	4
Medium Threat	177	6	1
Low Threat	21	58	30
Very Low Threat	26	182	43
OPERATOR Incident Ranking			
High Threat	13	30	8
Medium Threat	133	37	4
Low Threat	136	9	18
Very Low Threat	25	204	48

Table 3—Summary of Experimental Data Analyses

Row 4 presents the results from M-Correlator’s aggregation services. For example, from the 25,624 raw RealSecure alerts processed by M-Correlator, 250 correlated security incidents were created as security incidents. M-Correlator employed two alert aggregation clauses to perform the alert reduction, as discussed in Section Alert Aggregation Rules. Row 4 provides a further classification of alerts as follows:

- Privilege alert — activity that attempts to subvert administrative or user privileges
- Bad access alert — activity that attempts to violate access control policies through either a read or execute access violation, an integrity violation, user environment corruption, system environment corruption, or binary subversion, (e.g., a Trojan horse or virus infection)

- Availability alert — activity that appears to indicate a denial of service or computing asset in distress or unresponsive
- Probe alert — activity that matches intelligence gathering probes or sweeps of the network address space for systems or services
- Suspicious alerts — activity that is highly indicative of potential computer misuse and that may be important for forensic purposes

Rows 5 and 6 present the results of the M-Correlator incident ranking scheme as performed on the security administrator and availability operator mission specifications, respectively. Ranked security alerts were collectively binned under four ranges of importance. Under eAMI these ranges and number of folders used to represent M-Correlator results are completely user-definable. For the purposes of this experiment, the incident rankings were categorized as follows:

- high threat– security incidents whose rank exceeds 150
- medium threat – security incidents ranking from 140 to 150
- low threat – security incidents ranking from 128 to 139 (as discussed in Section *Incident Rank Calculation*, 128 represents a neutral priority score)
- very low threat – security incidents whose rank is below 128.

EXAMPLE – THE .69 NETWORK SECADMIN VIEW

M-Correlation provides a second layer of results refinement after alert aggregation. By ranking incidents by their likely impact on mission operations, M-Correlator provides the operator a further degree of automated results reduction, by focusing the operator on the critical security incidents that must be addressed first. For example, RealSecure produced a total of 25,624 security alerts for the .69 network over the collection period. M-Correlator reduced this raw alert volume down to 250 raw

security incidents, and then ranked these alerts against the security administrator’s mission specification. As a result, of the 250 security incidents produced by M-Correlator, only four alerts reached an incident ranking greater than 150. The High Threat folder view of the .69 network security administrator is illustrated in Figure 8.

Four buffer overflow incidents are listed in the High Threat folder. M-Correlator found that in addition to being of high interest to the security administrator’s interest profile, the alerts target critical servers. These findings are reflected in the Priority column of High Threat folder view in Figure 8. In addition, M-Correlator was also able to corroborate, by using its topology database, that the vulnerable services required for these attacks to succeed were indeed present on the target systems.

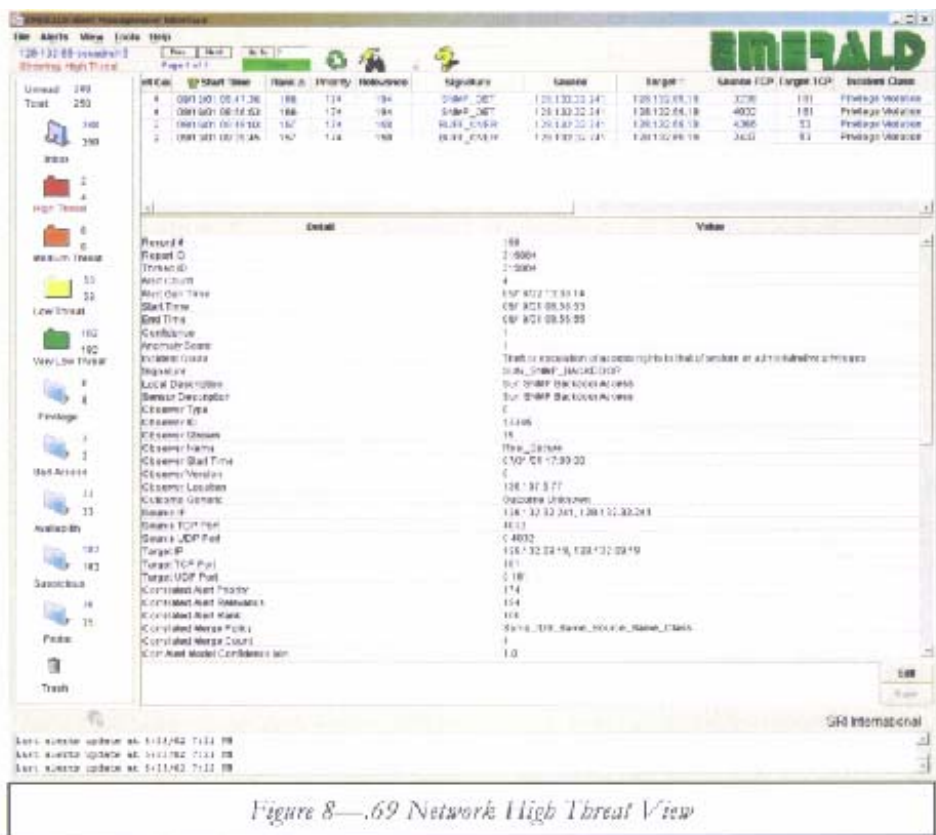


Figure 8—.69 Network High Threat View

Figure 9 illustrates the view of the opposite spectrum of incident ranking importance. This figure shows the Very Low Threat folder view for the security administrator of the .69 network. Alerts reviewed in this folder were found to be of low user interest, involve noncritical assets and network services, and involve vulnerabilities that could either not be validated against the asset's topology or explicitly found to be irrelevant to the target asset's configuration. In the example of the .69 network, roughly 70% of the security incidents analyzed by M-Correlator were found to be of very low interest to the security administrator, and could be set aside for later offline review or forensic analysis, as required.

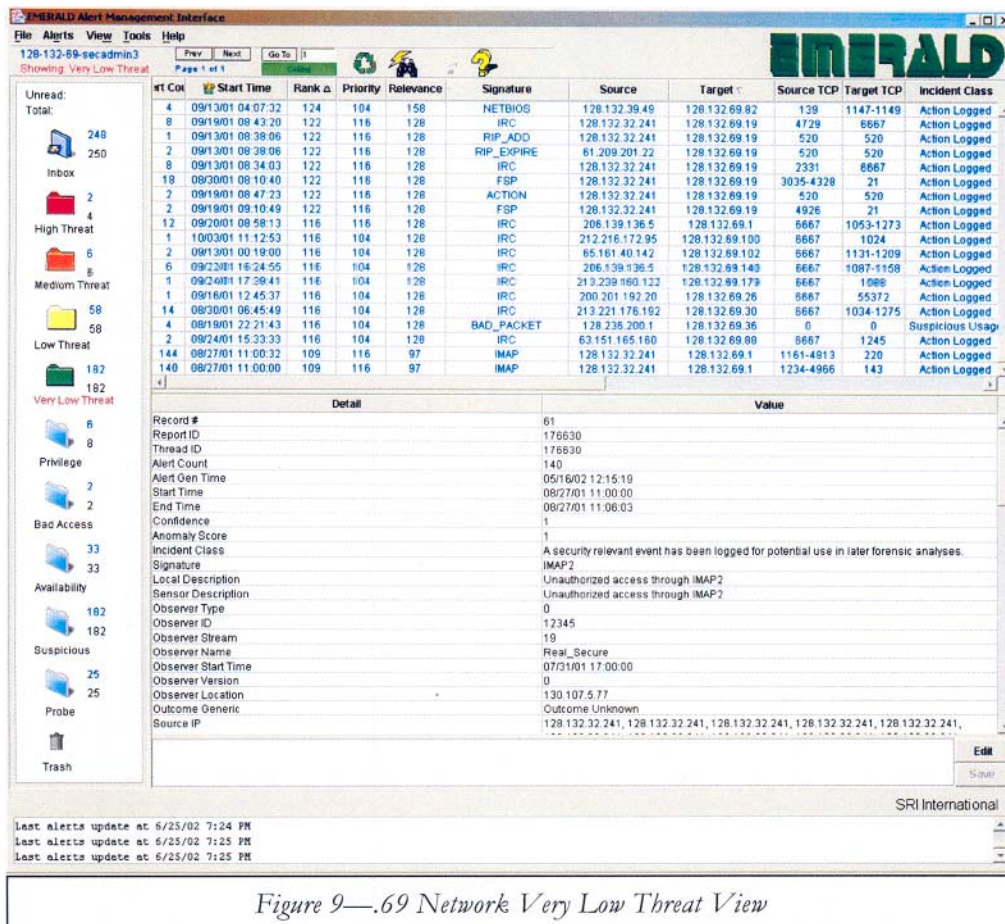


Figure 9—.69 Network Very Low Threat View

EXAMPLE – THE .69 NETWORK OPERATOR VIEW

To demonstrate the diversity of perspectives that may be supported within M-Correlator, SRI developed the mission specification of a network administrator whose responsibilities primarily involve the ensured availability of critical assets and services. This view of RealSecure can be supported simultaneously and could in fact be managed by the same RealSecure user, who wants to prioritize alerts from two perspectives.

While the operator and security administrator share a common view of which assets in the .69 network are critical, they have independent views of which incident types are of highest threat to their missions. In the case of the security administrator, privilege violations were of greatest importance, and were therefore elevated to the high rank, as shown in Figure 10. In the operator's case, maintaining the availability of critical servers, such as the two Windows assets 128.132.69.75 and 128.132.69.76 is of the highest importance. This is reflected in the operator's High Threat folder, as shown in Figure 10. This display illustrates that M-Correlator is observing a series of failures by other .mil systems in accessing NetBIOS services on the two critical hosts. RealSecure reports these failed access attempts as `WINDOWS_ACCESS_ERROR`, which are classified as asset distress incidents by the M-Correlator's alert normalization service. M-Correlator helps the operator focus in on those INFOSEC reports indicating availability problems to critical systems first, while still making available the breadth of security information for later review or forensic analysis.

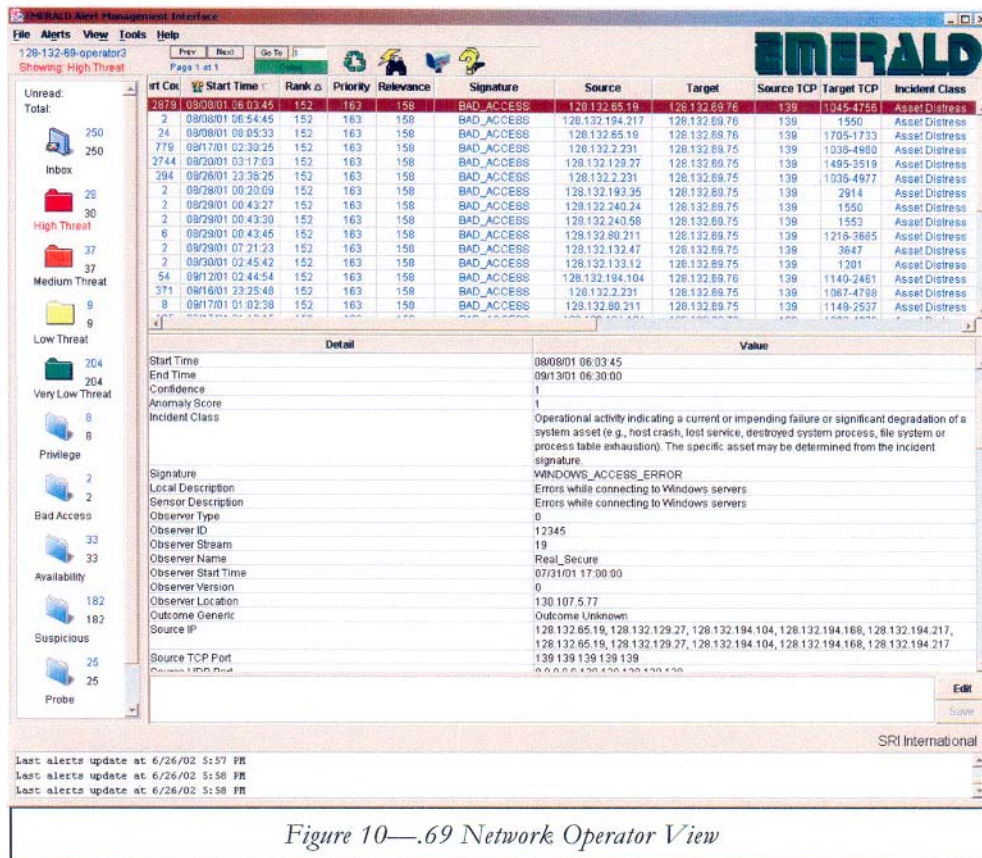


Figure 10—.69 Network Operator View

Observations and Lessons Learned

Conducting an experiment on the efficacy of M-Correlator on a large set of operational data proved to be a valuable learning experience, and in many ways produced results that were very consistent with SRI's own independent testing of this prototype system. The following is a brief enumeration of observations that are not only applicable to the AFRL AFED data analysis, but have been experienced in recent months as M-Correlator has been employed to examine ASIM, Snort, EMERALD, and other sensor data provided by third parties as well as live deployments within SRI internal networks.

UNINTERESTING SUBNETWORKS

Like several IDS sensors explored by SRI, the security alerts produced by ISS RealSecure regarding Rome Laboratory networks were categorized into two general classes of “uninterestingness”:

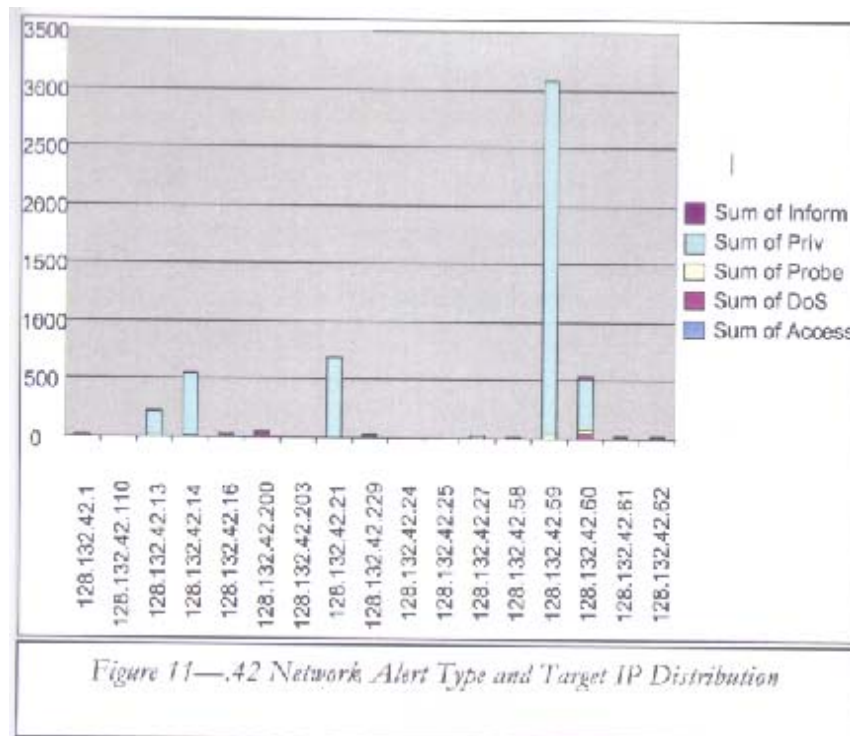
- A large number of ISS RealSecure alerts were related to internal LANs in which a single alert type was repeated tens of thousands of times for one or two hosts within the LAN. This was possibly due to a very strict firewall policy in which very few hosts are allowed to exchange internal-to-external communication, and methods for communication are highly restrictive. In addition, these situations have a high propensity to be caused by false positives or a mismatch in the sensor’s surveillance policy and the typical use of the exposed system (e.g., IRC communication alerts in environments where IRC is allowed, or legal NetBIOS communication between an internal host and other systems in the network).
- The entire IP space of a LAN is represented in sensor alerts regarding a single alert type. This is usually indicative of sweeps with a significant potential that the IP space of the LAN are phantom destinations.

THE POWER OF BASIC AGGREGATION

M-Correlator was highly successful in its ability to aggregate alerts. In all cases M-Correlator was able to reduce the raw alert count into security incident count by two or more orders of magnitude. This is not a unique experience to AFRL AFED ISS RealSecure reports, as SRI has experienced similar reductions in Snort and ASIM alerts on different operational networks. In reviewing these data sets, there appear to be a few basic properties of the alert streams that allow these levels of reduction. It should be noted that to date, the primary experiences in M-Correlator alert analysis

involve inward-looking intrusion detection sensors that are placed behind a firewall. The conditions described here may not be applicable to, for example, intrusion detection systems deployed at peering points (such as sensors deployed to monitor an ISP's client networks or a very large enterprise).

It appears that alert production is not, by nature, evenly distributed across all rules in the rule-base.

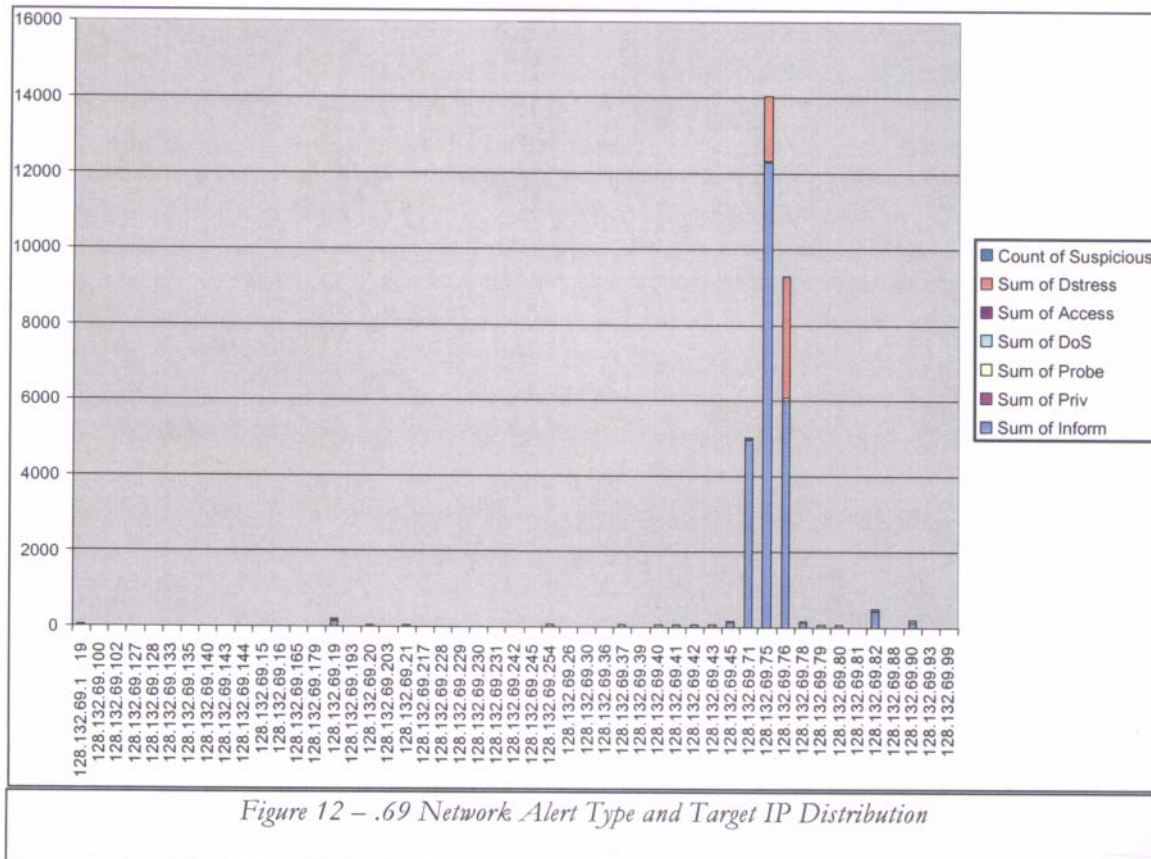


Rather than lack of diversity from external threat methods, this lack of alert diversity appears to occur with two common conditions. First, sensors placed behind firewalls are exposed to traffic that has been filtered to a very minimal set of exposed network interfaces, and thus a smaller set of vulnerabilities is available to attackers. Second, sensor heuristics themselves are not all implemented to provide the same level of fidelity. Either by accident or on purpose sensors will incorporate some

heuristics that are very prone to activate on nonmalicious traffic and other alerts that are highly targeted to specific traffic.

In addition, alert production from sensors behind a firewall tend to produce alerts that span a very small number of IP addresses in the total LAN IP space. Of course, this observation excludes situations such as those discussed in the previous section in which probes are allowed through the firewall and attempt to span the entire LAN IP space. The lack of IP space diversity enhances a correlation system's ability to fuse alerts to common targets.

To illustrate these points, consider the Target IP and Alert Type distributions found in the RealSecure alerts of the AFRL AFED 128.132.42.X and 128.132.69.X networks, illustrated in Figure 11 and Figure 12. The X-axis in each graph enumerates IP addresses with associated RealSecure alerts, and the Y-axis represents the total number of alerts. For each IP address, the graph illustrates the total number of alerts that targeted the IP address, as well as a breakdown of alerts by incident class. Both networks are shown to have a few IP addresses that receive the overwhelming majority of alerts, and alert diversity is very minimal. Very similar alert distributions have been observed in



other data set analyses performed by M-Correlator involving different target networks and IDS sensors.

INCIDENT RANKING SURVIVES SPARSE ALERT CONTENT

Among the greatest dependencies of M-Correlator with respect to the quality of its ranking calculation is the need for high-quality sensor input. Not all IDS sensors are capable of capturing the same amount of information per alert, and the alert content produced by ISS RealSecure is considered very minimal relative to the quality of content captured by other sensors such as ASIM, EMERALD, and Snort. With respect to ISS RealSecure, each alert was populated with the following content:

- Source IP address — indicates the source IP address of the packet(s) that caused a RealSecure heuristic to activate. While most RealSecure alerts represented have a corresponding relationship between the source address and attacker, this is not always the case, and RealSecure provides no indicator when it is not. (A similar inconsistency occurs with Snort. However, the EMERALD NIDS appliance version of Snort has been modified to resolve this inconsistency.)
- Destination IP address — indicates the destination of the packet(s) that caused a RealSecure heuristic to activate.
- Timestamp — indicates the sensor local time at which the alert was produced
- Source port — indicates the source port of the packet that activated the RealSecure heuristic. This field does not indicate the protocol (UDP or TCP) and is not provided on ICMP alerts.
- Target port — indicates the target port of the packet that activated the RealSecure heuristic. This field does not indicate the protocol (UDP or TCP) and is not provided on ICMP alerts.

RealSecure alerts provide M-Correlator an extremely minimal amount of information from which to perform incident ranking and alert aggregation. Nevertheless, even with this bare minimum of alert information, M-Correlator provided a meaningful separation in the alert rankings of the three analyzed networks, and reduced the total alert stream by two or more orders of magnitude. These results help to demonstrate that M-Correlator can be an effective tool, even when processing INFOSEC device alerts that contain a bare minimum of information.

FUTURE WORK

The AFRL AFED environment and expertise of AFED personnel have provided an important experimentation and assessment of the value of mission-based alert correlation on INFOSEC device alerts from a complex network. Feedback from AFRL AFED personnel has led to several potential adjustments and extensions to the M-Correlator rank formulation procedure, and has provided a strong vetting of the practical aspects of operating M-Correlator on low-content data sets.

In July 2002, the EMERALD team completed its data set analysis and visited AFRL AFED personnel to present a full review of the M-Correlator results. The review of the M-Correlator experimental results demonstrated a significant incident reduction capability, coupled with an effective alert ranking system that can automatically isolate the highest-threat security incidents in the experimental data set provided by AFRL. Based on these results, AFRL AFED personnel requested that SRI produce a statement of work to extend the initial contract to provide an M-Correlator release that is integrated into the AFRL AFED system.

Upon invitation by AFRL AFED personnel, SRI has submitted a proposal for a collaborative effort to embed a release of M-Correlator into the AFED/AFRL system. The objective is to deliver a permanent on-site M-Correlator release that can interact with operational AFED databases for the analysis of at least one AFRL AFED sensor.

SUMMARY

This Final Technical Report presents the setup, analysis results, and assessment of an operational test experiment of the SRI EMERALD M-Correlator system. The experiment was conducted with the cooperation of AFRL AFED personnel, and focused on the analysis INFOSEC device alerts in the context of the AFRL AFED system. The objective was to pragmatically assess M-Correlator's ability to process and provide meaningful new information from a large operational data set. In this case, the experimental data set consisted of a collection of more than 730,000 ISS RealSecure alerts produced over a two-month period within the Rome Laboratory network computing environment.

The M-Correlator analysis was conducted on a selection of three Rome Laboratory networks with large amounts of diverse alert content. With guidance from AFRL AFED personnel, SRI produced two hypothetical mission specifications: one mission specification from the perspective of a security administrator responsible for prevention of unauthorized access, and one mission specification representing a network operator whose main responsibility is to maintain the availability of critical computing assets.

The results of the experimentation demonstrated an ability by M-Correlator to reduce the raw RealSecure alert stream by more than two orders of magnitude, coupled with a very effective incident ranking system that accurately prioritized alerts with respect to their impact on the mission specification. These results were presented in detail to AFRL AFED personnel, who considered the results to represent a valuable contribution to their analytical capability. As a result, SRI has been invited to seek a follow-on project to collaborate with AFRL AFED personnel on the permanent integration of M-Correlator into the AFED technology suite.