

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188	
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 the collection of information. Send comment regarding this burden estimates 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 08/08/05		3. REPORT TYPE AND DATES COVERED Final progress report, 07/01/02 – 05/31/05
4. TITLE AND SUBTITLE Correlating Alerts Using Prerequisites of Intrusions: Towards Reducing False Alerts & Uncovering High Level Attack Strategies			5. FUNDING NUMBERS DAAD 19-02-1-0219	
6. AUTHOR(S) Peng Ning and Douglas S. Reeves				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) North Carolina State University, B Holladay Hall, Raleigh, NC 27695			8. PERFORMING ORGANIZATION	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U. S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSORING / MONITORING AGENCY REPORT NUMBER 43709.13-CI	
11. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.				
12 a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12 b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Traditional intrusion detection systems (IDSs) focus on low-level attacks or anomalies, and raise alerts independently, though there may be logical connections between them. In situations where there are intensive attacks, not only will actual alerts be mixed with false alerts, but the amount of alerts will also become unmanageable. As a result, it is difficult for human users or intrusion response systems to understand the alerts and take appropriate actions. The objective of this project is to develop techniques and tools to facilitate the automatic (or semi-automatic) analysis of IDS alerts. In particular, we have thoroughly investigated the following issues: construction of attack scenarios from IDS alerts, efficient and effective analysis of large sets of IDS alerts, learning of attack strategies from correlated alerts, hypothesizing and reasoning about attacks missed by IDSs, integration of intrusion evidence from IDSs and other complementary information sources, alert correlation when there are privacy concerns, systematic development of the knowledge base required for alert correlation in our approach, and vulnerability analysis of MANET routing protocols to facilitate the application of alert correlation in MANET applications. We have made significant progress in this project on all these issues, as described in the report.				
14. SUBJECT TERMS Security, Intrusion Detection, Alert Correlation, Wireless Networks			15. NUMBER OF PAGES 19	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OR REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION ON THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Final Progress Report

Table of Content

1	Statement of the problem studied	1
2	Summary of the most import results	2
2.1	<i>A formal framework for correlating intrusion alerts and constructing high-level attack scenarios</i>	<i>2</i>
2.2	<i>Interactive utilities for analyzing intensive intrusion alerts</i>	<i>2</i>
2.3	<i>Multi-level key chain method for scalable broadcast authentication in distributed sensor networks.....</i>	<i>3</i>
2.4	<i>Adapting query optimization techniques for efficient intrusion alert correlation</i>	<i>3</i>
2.5	<i>Insider attacks against mobile ad-hoc routing protocols</i>	<i>4</i>
2.6	<i>Learning attack strategies from intrusion alerts</i>	<i>4</i>
2.7	<i>Hypothesizing and reasoning about attacks missed by intrusion detection systems</i>	<i>4</i>
2.8	<i>Integrating complementary intrusion evidence from multiple information sources</i>	<i>5</i>
2.9	<i>Alert Correlation through Triggering Events and Common Resources</i>	<i>5</i>
2.10	<i>Privacy-Preserving Alert Correlation Using Concept Hierarchies</i>	<i>5</i>
3	List of all publications and technical reports supported under this grant.....	6
3.1	<i>Papers published in peer-reviewed journals</i>	<i>6</i>
3.2	<i>Papers published in non-peer-reviewed journals or in conference proceedings.....</i>	<i>6</i>
3.3	<i>Papers presented at meetings, but not published in conference proceedings</i>	<i>7</i>
3.4	<i>Manuscripts submitted, but not published</i>	<i>7</i>
4	List of all participating scientific personnel.....	7
5	Report of Inventions (by title only).....	7
6	Bibliography	7

1 Statement of the problem studied

Traditional intrusion detection systems (IDSs) focus on low-level attacks or anomalies, and raise alerts independently, though there may be logical connections between them. In situations where there are intensive attacks, not only will actual alerts be mixed with false alerts, but the amount of alerts will also become unmanageable. As a result, it is difficult for human users or intrusion response systems to understand the alerts and take appropriate actions.

The objective of this project is to develop techniques and tools to facilitate the automatic (or semi-automatic) analysis of IDS alerts. In particular, we have thoroughly investigated the following issues: (1) construction of attack scenarios from IDS alerts, (2) efficient and effective analysis of large sets of IDS alerts, (3) profiling and learning of attack strategies from correlated alerts, (4) hypothesizing and reasoning about attacks missed by IDSs, (5) integration of intrusion evidence from IDSs and other complementary information sources, (6) alert correlation when there are privacy concerns, (7) systematic development of the knowledge base required for

alert correlation in our approach, and (8) vulnerability analysis of MANET routing protocols to facilitate the application of alert correlation in MANET applications. We have made significant progress in this project on all these issues, as described next.

2 Summary of the most import results

Under the support of this grant, we have obtained a number of important results, which are summarized below:

2.1 A formal framework for correlating intrusion alerts and constructing high-level attack scenarios

Our first accomplishment is a formal framework for correlating intrusion alerts and constructing high-level attack scenarios based on the low-level intrusion alerts. Our techniques address an important problem in the management of IDSs. Our method can be explained easily based on the following observation: most intrusions are not isolated, but related as different stages of attacks, with the early stages preparing for the later ones. For example, in Distributed Denial of Service (DDOS) attacks, the attacker has to install the DDOS daemon programs in vulnerable hosts before he/she can instruct the daemons to launch an attack. In other words, an attacker has to (or usually does) reach a certain state before he/she can carry out certain attacks, and usually reaches the state by launching some other attacks.

Based on this observation, we correlate alerts using prerequisites and consequences of intrusions. Intuitively, the prerequisite of an intrusion is the necessary condition for the intrusion to be successful, while the consequence of an intrusion is the outcome of the intrusion if it is successful. For example, the existence of a vulnerable service is the prerequisite of a remote buffer overflow attack against the service, and as the consequence of the attack, the attacker may gain access to the host. Accordingly, we correlate the alerts together when the attackers launch some early attacks to prepare for the prerequisites of some later ones. For example, if they use a UDP port scan to discover the vulnerable services, followed by an attack against one of the services, we can correlate the corresponding alerts together. In addition, our formalism provides an intuitive representation of correlated alerts and a specific mechanism for alert correlation, which leads to our implementation of the method.

The details of the above result can be found in papers No. 4 and No. 12 listed in section 3.

2.2 Interactive utilities for analyzing intensive intrusion alerts

Our second accomplishment is the development of a set of utilities for interactive analysis of intensive intrusion alerts. Each utility takes a set of hyper-alerts as input. Depending on the output, these utilities can be divided into two classes: hyper-alert generating utilities and feature extraction utilities. A hyper-alert generating utility outputs one or multiple sets of hyper-alerts, while a feature extraction utility only outputs the properties of the input hyper-alerts. We have developed six utilities, including alert aggregation/disaggregation, focused analysis, clustering analysis, frequency analysis, link analysis, and association analysis. The first three utilities are hyper-alert generating utilities, while the last three are feature extraction utilities.

These utilities are intended for human users to analyze and understand the correlated alerts as well as the strategies behind them. We studied the effectiveness of these utilities through a case study with the network traffic captured at the DEF CON 8 Capture the Flag (CTF) event. Our results showed that they could effectively simplify the analysis of large amounts of alerts. Our analysis also revealed several attack strategies that appeared in the DEF CON 8 CTF contest.

The details of the development of the three utilities and the experimental evaluation of them can be found in papers No. 4 and No. 13 in section 3.

2.3 Multi-level key chain method for scalable broadcast authentication in distributed sensor networks

In studying the applicability of our techniques to wireless networks, we developed a multi-level key chain method based on μ TESLA to enable cost-effective and scalable broadcast authentication in low-end sensor networks. Broadcast authentication is an essential service in distributed sensor networks. Because of the large numbers of sensor nodes and the broadcast nature of the communication in distributed sensor networks, it is usually desirable for the base stations to broadcast commands and data to the sensor nodes. In hostile environments (e.g., battle field, anti-terrorists operations), it is necessary to enable the sensor nodes to authenticate the broadcast messages received from the base station

The only practical broadcast authentication protocol in low-end sensor networks (before our result) is μ TESLA, which is adapted from a stream authentication protocol called TESLA. However, μ TESLA requires that the base station unicast the initial parameters to the sensor nodes individually. This feature severely limits the application of μ TESLA in large sensor networks. To address this problem, we developed an extension to μ TESLA. The basic idea is to *predetermine* and *broadcast* the initial parameters required by μ TESLA instead of unicast-based message transmission. In the simplest form, our extension distributes the μ TESLA parameters during the initialization of the sensor nodes (e.g., along with the master key shared between each sensor and the base station). To provide more flexibility, especially to prolong the lifetime of μ TESLA without requiring a very long key chain, we introduced a multi-level key chain scheme, in which the higher-level key chains are used to authenticate the commitments of lower-level ones. To further improve the survivability of the scheme against message loss and Denial of Service (DOS) attacks, we used redundant message transmission and random selection strategies to deal with the messages that distribute key chain commitments. The resulting scheme removes the requirement of unicast-based initial communication between base station and sensor nodes while keeping the nice properties of μ TESLA (e.g., tolerance of message loss, resistance to replay attacks). Our implementation and experiments further demonstrated that our scheme could tolerate high channel loss rate and is resistant to certain known DOS attacks to a certain degree.

The details of the above result can be found in papers No. 3 and No. 11 listed in section 3.

2.4 Adapting query optimization techniques for efficient intrusion alert correlation

Intrusion alert correlation is the process to identify high-level attack scenarios by reasoning about low-level alerts raised by IDSs. The efficiency of intrusion alert correlation is critical in enabling interactive analysis of intrusion alerts as well as prompt responses to attacks. In this part of work, we adapted a number of main memory index structures (e.g., T Trees, Linear Hashing) and database query optimization techniques (e.g., nested loop join, sort join) to speed up intrusion alert correlation. By taking advantage of the characteristics of the alert correlation process, we developed three techniques named *hyper-alert container*, *two-level index*, and *sort correlation*. We also performed a series of experiments designed to evaluate the effectiveness of these techniques. These experiments demonstrate that (1) hyper-alert containers improve the efficiency of order-preserving index structures (e.g., T Trees), with which an insertion operation involves search, (2) two-level index improves the efficiency of all index structures, (3) a two-level index structure combining Chained Bucket Hashing and Linear Hashing is the most efficient for streamed alerts with and without memory constraint, and (4) sort correlation with heap sort algorithm is the most efficient for alert correlation in batch.

More details can be found in paper No. 9 listed in section 3.

2.5 Insider attacks against mobile ad-hoc routing protocols

In this part of work, we performed a systematic analysis of insider attacks against mobile ad-hoc routing protocols, using the Ad hoc On-Demand Distance Vector (AODV) protocol as an example. We identify a number of attack goals and then study how to achieve these goals through misuses of the routing messages. To facilitate the analysis, we classify the insider attacks into two categories: *atomic misuses* and *compound misuses*. Atomic misuses are performed by manipulating a single routing message, which cannot be further divided; compound misuses are composed of combinations of atomic misuses and possibly normal uses of the routing protocol. Our analysis reveals several classes of insider attacks, including route disruption, route invasion, node isolation, and resource consumption. To validate our results, we have implemented and studied these attacks through simulation, which demonstrate the impact of these attacks.

Additional details can be found in papers No. 1 and No. 10 listed in section 3.

2.6 Learning attack strategies from intrusion alerts

Understanding the strategies of attacks is crucial for security applications such as computer and network forensics, intrusion response, and prevention of future attacks. In this part of research, we developed techniques to automatically learn attack strategies from intrusion alerts. Central to these techniques is a model that represents an attack strategy as a graph of attacks with constraints on the attack attributes and the temporal order among these attacks. To learn the intrusion strategy is then to extract such a graph from a sequence of intrusion alerts. To further facilitate the analysis of attack strategies, which is essential to many security applications such as computer and network forensics and incident handling, we developed techniques to measure the similarity between attack strategies. The basic idea is to reduce the similarity measurement of attack strategies into error-tolerant graph isomorphism problem, and measure the similarity between attack strategies in terms of the cost to transform one strategy into another. Finally, we have performed a series of experiments, which demonstrate the potential of the aforementioned techniques.

Additional details about this work can be found in paper No. 8 listed section 3.

2.7 Hypothesizing and reasoning about attacks missed by intrusion detection systems

Several alert correlation methods were proposed in the past several years to construct high-level attack scenarios from low-level intrusion alerts reported by IDSs. These correlation methods have different strengths and limitations; none of them clearly dominate the others. However, all of these methods depend heavily on the underlying IDSs, and perform poorly when the IDSs miss critical attacks. In order to improve the performance of intrusion alert correlation and reduce the impact of missed attacks, we developed a series of techniques to integrate two complementary types of alert correlation methods: (1) those based on the similarity between alert attributes, and (2) those based on prerequisites and consequences of attacks. In particular, we developed techniques to hypothesize and reason about attacks possibly missed by IDSs based on the indirect causal relationship between intrusion alerts and the constraints they must satisfy. We also investigated additional techniques to validate the hypothesized attacks through raw audit data and to consolidate the hypothesized attacks to generate concise attack scenarios. Our experimental results demonstrate the potential of these techniques in building high-level attack scenarios and reasoning about possibly missed attacks.

Additional details about this work can be found in paper No. 2 and No. 7 listed in section 3.

2.8 Integrating complementary intrusion evidence from multiple information sources

We developed techniques to integrate and reason about complementary intrusion evidence such as intrusion alerts generated by IDSs and reports by system monitoring or vulnerability scanning tools. To facilitate the modeling of intrusion evidence, we classify intrusion evidence into either *event-based evidence* or *state-based evidence*. Event-based evidence refers to observations (or detections) of intrusive actions (e.g., IDS alerts), while state-based evidence refers to observations of the effects of intrusions on system states. Based on the interdependency between event-based and state-based evidence, we developed techniques to automatically integrate complementary evidence into Bayesian networks, and reason about uncertain or unknown intrusion evidence based on verified evidence. The experimental results demonstrated the potential of the proposed techniques. In particular, additional observations by system monitoring or vulnerability scanning tools can potentially reduce the false alert rate and increase the confidence in alerts corresponding to successful attacks.

Additional details about this work can be found in paper No. 5 listed in section 3.

2.9 Alert Correlation through Triggering Events and Common Resources

Complementary security systems are widely deployed in networks to better protect digital assets. Alert correlation is essential to understand the security threats and take appropriate actions. Here we propose a novel correlation approach based on triggering events and common resources. One of the key concepts in our approach is triggering events, which are the (low-level) events that trigger alerts. By grouping alerts that share "similar" triggering events, a set of alerts can be partitioned into different clusters where each cluster may correspond to the same attack. Our approach further examines whether the alerts in each cluster are consistent with relevant network and host configurations, which helps analysts partially identify the severity of alerts and clusters. The other key concept in our approach is input and output resources. Intuitively, input resources are the necessary resources for an attack to succeed, and output resources are the resources that an attack supplies if successful. We model each attack through specifying input and output resources. By identifying the "common" resources between output resources of one attack and input resources of another, our approach discovers causal relationships between alert clusters and builds attack scenarios. The preliminary experimental results demonstrate the usefulness of the proposed techniques.

We also investigated how to use resources to facilitate the specification of hyper-alert types (or, equivalently, attacks), which are required in our correlation method. We organize resources into trees, where the nodes in the trees are labeled with conditions (represented by predicates). To specify the prerequisite and consequence of an attack, we first look for the desirable resource trees related to the attack's prerequisite and consequence, then traverse the trees to find the appropriate nodes, and finally select the suitable predicates to put into the prerequisite and consequence. Our approach is simple and less expert-dependent. The usability study and comprehensiveness study (with more than 350 attack types) demonstrate the effectiveness of our approach.

Additional details about this work can be found in papers No. 6 and No. 16 listed in section 3.

2.10 Privacy-Preserving Alert Correlation Using Concept Hierarchies

With the increasing security threats from infrastructure attacks such as worms and distributed denial of service attacks, it is clear that the cooperation among different organizations is necessary to defend against these attacks. However, organizations' privacy concerns for the incident and security alert data require that sensitive data be sanitized before they are shared with other organizations. Such sanitization process usually has negative impacts on intrusion analysis (such as alert correlation). To balance the privacy requirements and the need for intrusion analysis, we propose a privacy-preserving alert correlation approach based on concept hierarchies. Our approach consists of two phases. The first phase is *entropy guided alert sanitization*, where sensitive alert

attributes are generalized to high-level concepts to introduce uncertainty into the dataset with partial semantics. To balance the privacy and the usability of alert data, we propose to guide the alert sanitization process with the entropy or differential entropy of sanitized attributes. The second phase is *sanitized alert correlation*. We focus on defining similarity functions between sanitized attributes and building attack scenarios from sanitized alerts. Our preliminary experimental results demonstrate the effectiveness of the proposed techniques in terms of various measures (e.g., correct classification rates, false alert rates, and detection rates).

More details of this work can be found in paper No. 15 in section 3.

3 List of all publications and technical reports supported under this grant

3.1 Papers published in peer-reviewed journals

1. Peng Ning, Kun Sun, "How to Misuse AODV: A Case Study of Insider Attacks against Mobile Ad-hoc Routing Protocols," (Extended version of Paper No. 10) To appear in *Ad Hoc Networks Journal*. 2005.
2. Peng Ning, Dingbang Xu, "Hypothesizing and Reasoning about Attacks Missed by Intrusion Detection Systems," in *ACM Transactions on Information and System Security*, Volume 7, Number 4, pages 591--627, November 2004.
3. Donggang Liu, Peng Ning, "Multi-Level μ TESLA: Broadcast Authentication for Distributed Sensor Networks," *ACM Transactions in Embedded Computing Systems (TECS)*, Vol. 3, No. 4, pages 800--836, November 2004.
4. Peng Ning, Yun Cui, Douglas S. Reeves, and Dingbang Xu, "Tools and Techniques for Analyzing Intrusion Alerts," in *ACM Transactions on Information and System Security*, Volume 7, Number 2, pages 273--318, May 2004.

3.2 Papers published in non-peer-reviewed journals or in conference proceedings

5. Yan Zhai, Peng Ning, Purush Iyer, Douglas S. Reeves, "Reasoning about Complementary Intrusion Evidence," in *Proceedings of 20th Annual Computer Security Applications Conference*, pages 39--48, December 2004.
6. Dingbang Xu, Peng Ning, "Alert Correlation through Triggering Events and Common Resources," in *Proceedings of 20th Annual Computer Security Applications Conference*, pages 360--369, December 2004.
7. Peng Ning, Dingbang Xu, Christopher G. Healey, and Robert A. St. Amant, "Building Attack Scenarios through Integration of Complementary Alert Correlation Methods," in *Proceedings of the 11th Annual Network and Distributed System Security Symposium (NDSS '04)*, pages 97--111, February 2004.
8. Peng Ning, Dingbang Xu, "Learning Attack Strategies from Intrusion Alerts," in *Proceedings of the 10th ACM Conference on Computer and Communications Security (CCS '03)*, pages 200--209, Washington D.C., October 2003.
9. Peng Ning, Dingbang Xu, "Adapting Query Optimization Techniques for Efficient Intrusion Alert Correlation," in *Proceedings of the 17th IFIP WG 11.3 Working Conference on Data and Application Security*, August 2003.
10. Peng Ning, Kun Sun, "How to Misuse AODV: A Case Study of Insider Attacks against Mobile Ad-hoc Routing Protocols," in *Proceedings of the 4th Annual IEEE Information Assurance Workshop*, pages 60-67, West Point, June 2003. (**Best Paper Award**)
11. Donggang Liu, Peng Ning, "Efficient Distribution of Key Chain Commitments for Broadcast Authentication in Distributed Sensor Networks," in *Proceedings of the 10th Annual Network and Distributed System Security Symposium*, pages 263--276, February 2003.
12. Peng Ning, Yun Cui, Douglas S. Reeves, "Constructing Attack Scenarios through Correlation of Intrusion Alerts," in *Proceedings of the 9th ACM Conference on Computer & Communications Security*, pages 245--254, Washington D.C., November 2002.

13. Peng Ning, Yun Cui, Douglas S. Reeves, "Analyzing Intensive Intrusion Alerts Via Correlation," in *Proceedings of the 5th International Symposium on Recent Advances in Intrusion Detection (RAID 2002)*, LNCS 2516, pages 74--94, Zurich, Switzerland, October 2002.

3.3 Papers presented at meetings, but not published in conference proceedings

14. Peng Ning, Yun Cui, Douglas S. Reeves, and Dingbang Xu, "Towards Automating Intrusion Alert Analysis," in *2003 Workshop on Statistical and Machine Learning Techniques in Computer Intrusion Detection*, September 2003.

3.4 Manuscripts submitted, but not published

15. Dingbang Xu, Peng Ning, "Privacy-Preserving Alert Correlation: A Concept Hierarchy Based Approach," Submitted for conference publication, June 2005.
16. Jaideep Mahalati, Dingbang Xu, Peng Ning, "Facilitating Alert Correlation Using Resource Trees," Submitted for conference publication, June 2005.
17. Yan Zhai, Peng Ning, Jun Xu, "Integrating IDS Alert Correlation and OS-Level Dependency Tracking," Submitted for conference publication, July 2005.

4 List of all participating scientific personnel

- Dr. Peng Ning (PI)
- Dr. Douglas S. Reeves (Co-PI)
- Yun Cui, MS in Computer Science, degree earned in December 2002
- Alfredo Serrano, MS in Computer Science, degree earned in May 2003
- Yiquan Hu, MS in Computer Science, degree earned in December 2003
- Donggang Liu, PhD in Computer Science, degree earned after the employment on the project
- Kun Sun, PhD in progress
- Yan Zhai, PhD in progress
- Pan Wang, PhD in progress

5 Report of Inventions (by title only)

- A method to correlate intrusion alerts based on prerequisites and consequences of attacks
- Utilities for interactively analyzing intrusion alerts
- Adapting query optimization techniques for efficient intrusion alert correlation
- Learning attack strategies from intrusion alerts
- Hypothesizing and reasoning about attacks missed by intrusion detection systems
- Integrating complementary intrusion evidence from multiple information sources
- Alert correlation through triggering events and common resources
- Privacy-preservation alert correlation using concept hierarchies
- Efficient distribution of key chain commitments for broadcast authentication in wireless sensor networks

6 Bibliography

N/A