

Forensic Analysis of Risks in Enterprise Systems

Peter Stephenson, CISSP, CISM, FICAF

Research Scientist

Center for Regional and National Security, Eastern Michigan University

Abstract: *The concepts of vulnerability assessment and penetration testing as methods of risk analysis have been a staple of the practice of information security. The seminal paper by Farmer and Venema [FV92] introduced the concept of performing penetration tests as a method of vulnerability assessment. Since the early 1990s the practices of vulnerability assessment and risk analysis have alternately converged and diverged as new methods waxed and waned. Most recently standards such as BS7799/ISO17799 have focused on the synergy between technical testing and risk analysis.*

Implicit in the concepts of vulnerability analysis and risk analysis, whatever the techniques used, is the notion of risk management. Risk management returns to first principles of vulnerabilities, threats, impacts and countermeasures. Current thinking, as embodied in various standards and industry-specific regulations, implies a holistic approach to risk management that comprises technical, operational, and administrative controls and the required assessments to establish their efficacy in managing the organization's information technology (IT) risks.

Unfortunately, there is, today, limited evidence of credible risk analysis procedures that combine technical and non-technical assessment and analysis methods satisfactorily. Risk analysis methods tend to focus on expected annual loss while vulnerability and penetration testing methods focus on uncovering holes in the system that would permit an intruder to compromise the organization's information assets. There are serious limitations to both of those approaches.

The results of various methods of risk analysis in common use in IT systems today are suspect because the source data and assumptions upon which their conclusions are based is subjective and may be flawed and inconsistent. The results of various vulnerability and penetration assessments are of limited value because they are based upon the assumption that all possible vulnerabilities can be known and tested for, clearly not a possibility.

This paper discusses a novel technique that manages risks to an enterprise in terms of how well hardened, technically, operationally and administratively, the enterprise is against attack. Using the concept of formal analysis of covert channels between security policy domains, the Forensic Analysis of Risks in Enterprise Systems (FARES) process addresses threats, vulnerabilities, impacts and countermeasures from the perspective of forensic analysis of target enterprises responding to various threat models. Additionally, using a framework based upon the BS7799 standard, FARES can address all non-technical aspects the enterprise as those aspects might impact risk.

Because FARES also addresses the need for countermeasures, the output of a FARES analysis has the benefit of suggesting appropriate countermeasures and highlights countermeasures that while expensive and complex to implement, may not offer acceptable pay-back in terms of real protection. Other benefits of the FARES approach include identification of threats to as yet unidentified vulnerabilities, ease of regular updating over time, lower cost of performing the initial analysis compared with other forms of assessment and analysis, and the ability to re-assess quickly, easily and relatively inexpensively when the organization implements changes to the enterprise.

BACKGROUND

In 1995 the writer developed an approach to vulnerability assessment called VAST (Vulnerability Assessment SWAT Team) based upon the principle that the use of multiple assessment tools to gather large amounts of data about the enterprise would result in an acceptable picture of the enterprise's overall vulnerability.

Over time the notion of penetration testing to validate or invalidate discovered vulnerabilities led to a more focused approach and yielded more consistent and credible results. Statistical analysis of over 100 VAST engagements revealed an acceptable level of consistency.

However, over time it has become clear that the nature of hostile attacks has changed significantly enough that any type of assessment that depends for its success upon identifying known vulnerabilities only is flawed. Thus, given the rapidity with which new threats and vulnerabilities appear, and the emergence of so-called "zero-day" exploits¹, the VAST approach, however successful in the past, is underpowered for the future.

Since, today, virtually all vulnerability and penetration testing approaches rely on methods similar to those used in VAST, vulnerability and penetration testing are unacceptable for use on modern enterprises. The major weakness, from an information assurance perspective, is the inability to predict and test for unknown vulnerabilities. Virtually no approaches use the concept of threat modeling.

Additionally, by their nature, these approaches do not suggest appropriate countermeasures in most cases. One might characterize the testing approach as well as the approach to selecting countermeasures as "brute force".

An additional driver in today's organizational environments is the legal and regulatory climate in which the organization operates. Mapping the organization's risk to a set of standards, laws and/or regulations has become a necessity.

Current approaches attempt to satisfy this requirement by tailoring testing and analysis to

¹ Exploits that appear at the same time or nearly the same time as the discovery of the vulnerability they address.

those laws and regulations impacting the organization instead of hardening the enterprise and mapping the results of that hardening back to applicable laws, regulations and standards.

Current focus often is on incomplete "micro-testing" instead of more satisfactory holistic enterprise-wide risk management.

Risk Management Requirements

Classically, risk management addresses risk assessment, risk mitigation, security management and security auditing [RJ02]. The elements of a risk, as seen generally in quantitative risk assessments are threats, vulnerabilities, impacts and countermeasures (or "safeguards") [AS91].

Most risk analysis methods take these elements into account but, generally, do so in a manner that may be unreliable due to the reliability of source data. For example, estimating the impact of a particular type of attack against a particular environment is difficult. In the absence of hard historical data the estimation is subjective at best. Since most risk assessment methods require this type of data, it is reasonable to expect the results of the assessment to be no more reliable than the source data itself.

The "Bigger Hammer" Approach

A key component of risk analysis is vulnerability analysis (VA). VA today generally is performed either by scanning the target network with vulnerability scanners or by performing penetration testing. Both approaches require a comprehensive knowledge of vulnerabilities to be credible.

Because new vulnerabilities appear regularly - one source reported an average of 30 new vulnerabilities per week in the 12 months of April 2001 through March 2002 [ST02] - it is essentially impossible for testing tool developers to keep their tools current with the latest exploits. Because zero day exploits are becoming increasingly common that difficulty is magnified significantly.

It is equally unlikely that penetration teams will be able to keep up with the emergence of new vulnerabilities. Thus, dependence upon traditional vulnerability or penetration testing techniques is an ineffective way to supply the

vulnerabilities element of a risk assessment. Simply applying the “bigger hammer” of more vulnerability signatures for brute-force testing is not a rigorous approach to a meaningful understanding of vulnerabilities within an enterprise.

Threat Analysis

The second element of risk assessment is threat analysis. Threat profiling in the information protection domain is a badly misunderstood concept. Assessment of threat differs depending upon the individual, organization or agency performing the assessment

Jones [AJ02] defines risk as:

“The likelihood that a threat agent will successfully exploit a vulnerability to create an unwanted or adverse impact.”

He goes on to list the factors that must be considered when conducting a risk assessment as:

- The agent causing the threat
- The exploitable vulnerability
- The impact of a successful attack
- Mitigating factors

The concept of threats, and threat agents, therefore, is central to the conduct of a credible risk management program. Threats consist of natural threats, malicious threats and system issues.

In the case of natural threats we are, simply, concerned with identifying such threats and calculating the probability of occurrence. This is fairly straightforward since such natural events as earthquakes, storms and the like are the subject of numerous statistical analyses.

Malicious threat factors include:

- Capability
- Motivation
- Access
- Catalysts
- Inhibitors
- Amplifiers

These factors must be examined for each threat. By examining threats in detail we reduce dependency upon a foreknowledge of explicit

vulnerabilities since we can use broader vulnerability categories and determine, (a) if the category exists, and (b) if there is a credible threat/threat agent that could/would exploit it in some way. What the specific way is does not concern us. We are only concerned that a vulnerability category, such as denial of service, could be exploited and that there is a credible threat agent to attempt the exploit.

Impact Analysis

The notion of impact is the fuzziest of the elements making up risk. In virtually any organization, determining the impact of an event is a very subjective process usually resulting in a wide disparity of results from an equally wide community of assessors.

There is no easy way to include impact credibly in a risk assessment. The best that can be hoped for is a useful combination of historical records and consistency of analysis.

However, the more granularity with which the assessor attempts to measure impact, the more accurate, potentially is the result. The other side of that, however, is that it is not practical simply to add a large set of very granular impacts together to calculate an overall impact.

Impacts may impact each other as well as the enterprise as a whole. Thus, depending upon the nature of the enterprise, the overall impact of an event may be calculated simply, or it may be extremely complex with some impacts being cumulative. There may or may not be consistency in the cumulative nature of groups of impacts upon the whole.

Managing Risk

Returning to classical risk management requirements, the notions of security management and security audit are important. Security audit usually is considered to be useful for verifying that risk mitigation has been implemented and is effective.

However, this is a tedious process and often consists of annual, or less frequent, audits of information systems and system infrastructure. In very large enterprises, such audits often are conducted in segments over a period of years. Thus, a particular aspect of the enterprise may not get revisited by auditors with enough

frequency to be representative of the ongoing state of the enterprise.

Additionally, preparation for such an audit generally consists of vulnerability assessment followed by attempts at vulnerability mitigation to satisfy the auditors. Clearly, this approach is not in the best interests of the organization. However, it is not unexpected given the complex nature of large enterprises, the fast moving requirements for enterprise change to keep up with business needs and the paucity of skilled assessors.

Managing risk, like building security into applications, is best done as early in the applicable life cycle as possible. Generally, identifying and mitigating risks as the enterprise changes is easier and less expensive than global testing and mitigation just prior to an audit.

However, the cost of such “just in time” mitigation may still be excessive given that a detailed analysis and vulnerability test may need to be carried out, often by an expensive third party.

Thus, it is clear that some other method of risk management, one that can be carried out economically on a just in time basis as changes to the enterprise are introduced, would be beneficial, especially to organizations with very large enterprises.

It also is clear that depending upon risk management being “vulnerability-centric” is not practical. It is imperative that organizations manage risks as a whole, not exclusively the individual components of vulnerability, threat and impact. An effective risk management approach is likely to be holistic rather than specific to one area of risk.

A NEXT-GENERATION RISK MANAGEMENT PARADIGM

Current generation risk management approaches depend, to be successful, upon identifying and understanding the enterprise’s vulnerabilities. As we have pointed out that is a nearly impossible task in today’s computing environments. Therefore, a risk management protocol that does not suffer from that dependency may prove to be beneficial. Such an approach must take into consideration threats as well as vulnerabilities, the need for constant, and relatively simple,

updating of the enterprise’s risk profile and the cost-effective application of appropriate safeguards or countermeasures.

In suggesting such a next-generation approach, to be credible, we must begin from as little of the current generation as possible while not departing from first principles. Thus, we should avoid using techniques in current vogue adding only refinements. We should, however, retain as much of the underpinnings of risk management as necessary to maintain the integrity of the risk management process.

For our purposes we begin from the fundamental concepts suggested by Jones and others. We will use a traditional definition of risks, extending somewhat, but remaining consistent with Jones:

Risk consists of vulnerabilities, the credible threats that could exploit them in a particular environment, the impacts caused by such a successful exploitation, and the countermeasures used to mitigate the impact.

Qualitative vs. Quantitative Methods

The notion of performing a quantitative risk analysis depends upon the credibility of the input data. Often this data is very subjective. Such subjectivity, in turn, impacts the credibility of the results of the analysis.

Qualitative approaches are little better since even a qualitative analysis demands a set of decisions on the part of the analyst. Those decisions, often, are as subjective as the numbers hypothesized in quantitative assessments.

A next generation approach to risk management should make an attempt at avoiding the subjective limitations of current generation techniques. In terms of threats, vulnerabilities and countermeasures, costs, where they exist at all, are fairly straightforward to calculate. Difficulty in quantifying impacts, however, is subject to the same constraints in a next generation risk management approach as they are in current generation methods.

For those readers who may be interested, Figure 2 shows the architecture of the Knowledge Base (for source see footnote 2).

To the thirty general threats from the Knowledge Base, we may add those specific threats that

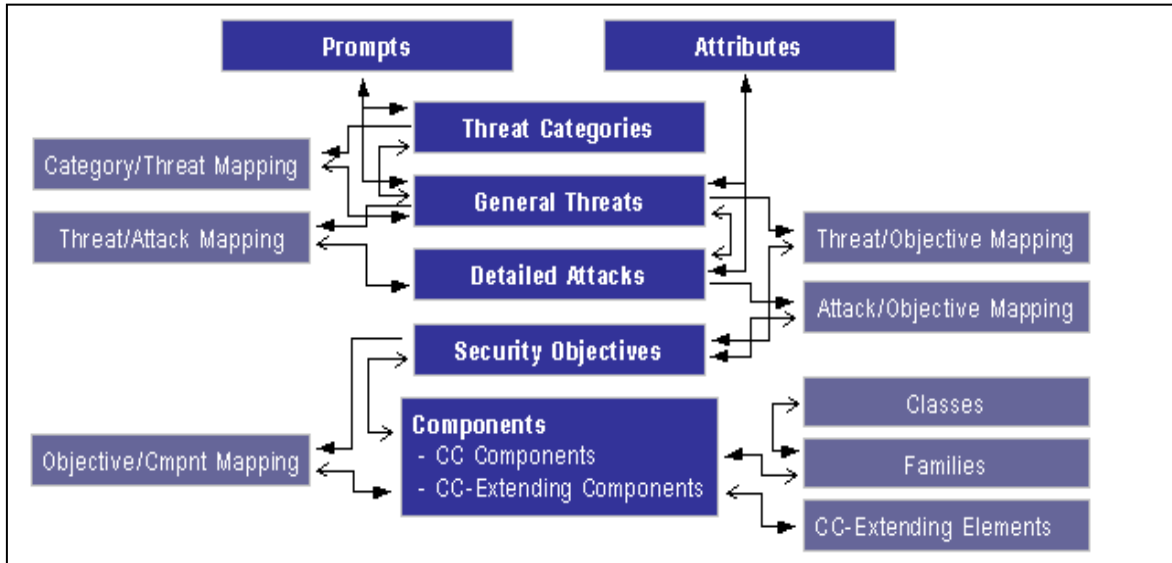


Figure 2 – Common Criteria Profiling Knowledge Base Architecture

A listing of general threats from the Knowledge base is in Table 1.

1.	Administrative errors of commission
2.	Administrative errors of omission
3.	Hostile administrator modification of user or system data
4.	Administrator violates user privacy policy
5.	A critical system component fails
6.	Software containing security-related flaws
7.	Failure of a distributed system component
8.	Hacker undetected system access
9.	Hacker attempts resource denial of service
10.	Hacker eavesdrops on user data communications
11.	Cryptanalysis for theft of information
12.	Hacker masquerading as a legitimate user or as system process
13.	Message content modification
14.	Exploitation of vulnerabilities in the physical environment of the system
15.	Social engineering
16.	Malicious code exploitation
17.	Unexpected disruption of system or component power
18.	Recipient denies receiving information
19.	Sender denies sending information
20.	A participant denies performing a transaction
21.	Legitimate system services are spoofed
22.	Hostile user acts cause confidentiality breaches
23.	User abuses authorization to collect data
24.	User errors cause confidentiality breaches
25.	User error makes data inaccessible
26.	User errors cause integrity breaches
27.	User errors undermine the system's security features
28.	User's misuse causes denial of service
29.	User abuses authorization to modify data
30.	User abuses authorization to send data

Table 1 – CC-PKB General Threats

relate explicitly to the domain under analysis.

By considering the threats against a security policy domain, we may extrapolate the vulnerabilities required for the threat to cause an impact. Knowing the potential threats and the vulnerabilities they address, we may analyze potential impact and craft appropriate countermeasures. We apply the threats to vulnerabilities after modeling the domain under test using CPNets.

To use the CC-PKB (as in Table 1) we simply apply each of the general threats to the security policy domains we determine for the enterprise under test. As we apply each general threat we consider Jones' framework for malicious threat factors. If the analysis of the general threat in light of Jones' factors yields a credible threat we consider it in our risk modeling. Otherwise we reject the general threat and move on to the next entry in the CC-PKB table, repeating the process as we go.

The notion of hypothesizing a threat/vulnerability/impact model and applying it to a CPNet of the domain under test is the theoretical equivalent of modeling the outcome of an actual incident with the exception that the incident in the theoretical case is a hypothetical model instead of an actual incident. By applying

countermeasures to the theoretical model such that the model deadlocks under threat, we identify cost-effective countermeasures against that threat.

Finally, we must analyze the generalized threats in accordance with Jones' approach to threat modeling. Not every possible general threat is credible in the context of a given policy domain under evaluation.

This approach focuses upon credible threats against an enterprise instead of attempting to identify all possible and potentially possible vulnerabilities within the enterprise and its individual components. Thus, it forms a manageable approach to the risk management of large enterprises. As we have seen, the notion of managing risks instead of managing vulnerabilities is more efficient and allows for response to unknown or potential vulnerabilities by identifying general threats that could exploit them.

required to perform this type of updating is considerably less than the cost for repeat vulnerability or penetration tests.

THE FARES PROCESS

There is a specific process for Forensic Analysis of Risks in Enterprise Systems. In simple terms it follows the process flow shown in Figure 3.

The process flow in Figure 3 actually is iterative in that it must be repeated for each policy domain and the interactions between domains. In a complex network this can be quite tedious the first time. However, understanding the network under test, its policy domains and testing each component and network segment using traditional tools is far more time consuming and, therefore, more expensive.

The FARES approach, however, is risk-centric instead of vulnerability-centric so it is reasonable to expect a significantly greater return on the

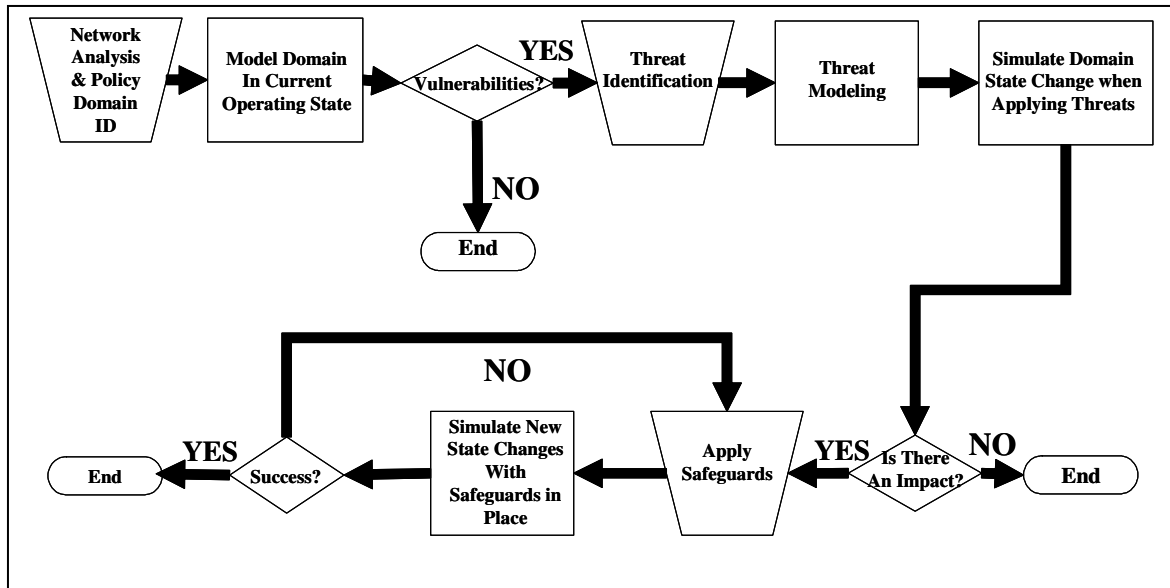


Figure 3 – FARES General Process Flow

A final benefit of such a next-generation approach to risk management is that it allows for rapid updating in the event of system changes. This updating requires simply that the altered portion of the policy domain under review be remodeled consistent with the updates or changes to it since the last model was created. The amount of time, resources and money

investment of time than one would expect from a vulnerability assessment that needed to be plugged into a risk analysis model.

The allocation of resources for a FARES analysis yields a risk management plan while the larger allocation of resources for a vulnerability assessment/penetration test yields a partial (at

best) picture of vulnerabilities only. Laboratory tests of the FARES approach as compared with the writer's experience with VAST vulnerability testing coupled with penetration tests has supported this contention.

Policy Domain Identification

The notion of policy domains is not new. The CORBA (Common Object Request Broker Architecture) Glossary [FS03] defines it as:

“A domain whose objects are all governed by the same security policy. There are several types of security policy domain, including access control policy domains.”

Smith, in a presentation for NIST [WS98] defines a security policy domain as:

“The scope over which a security policy is enforced. There may be subdomains for different aspects of this policy.”

Sanchez, Waitzman, Condell [SW00] et al describe security policy domains as:

“...an environment or context that is defined by a security policy, a security model, or architecture, and includes a set of system resources and a set of entities that have the right to access the resources.”

Regardless of the source of the definition, the concept is consistent: security policy domains are those logical and physical components of an enterprise governed by a single security policy. The policy may be explicit, as in a corporate policy, procedure or guideline, or, more commonly, implicit as in the configuration of devices governed by a policy. Typically, for practical purposes, we consider those configurations as the instantiation of the policy.

When identifying security policy domains we are concerned with both the scope of the domain and the interconnections between it and other security policy domains. These connections represent communications channels and they may be overt or covert. Additionally, we are concerned about covert storage channels that may exist within the security policy domain.

By basing our model upon a security policy

domain instead of individual components we both simplify the modeling process and focus upon hardening the enterprise at the domain level. While a security policy domain usually contains multiple components, it is not uncommon for a single, high sensitivity or high criticality component to comprise a single domain.

Current State Modeling and Simulation

The current operating state of each security policy domain is modeled based upon the domain's process and/or data flow. Since CPNets allow the creation of sub-nets, we can look at the overall enterprise as a single net comprising sub-nets, each of which represents a security policy domain.

Using Design/CPN or CPNTools³, the model can be run as a simulation once completed verifying that the model behaves as the security policy domain's current configuration dictates. It is important that the current state model be a correct representation of the actual operating state of the security policy domain it represents.

Once the individual security policy domain models are completed and verified, they can be connected to a super page that represents the enterprise as a whole or managed independently. Because security policy domains interact with those other domains to which they interconnect, it is important to simulate that interaction, however.

Threat Identification and Analysis

The threat identification and analysis process begins with the Profiling Knowledge Base. Each general threat in the knowledge base should be examined for applicability to each security policy domain in accordance with the procedure suggested by Jones [AJ02]. Once general threats have been identified and analyzed, special threats specific to the security policy domain in question should be identified and analyzed in the same manner. Specific threats may be identified through interviews with knowledgeable individuals having specific experience with the security policy domain in question.

Not all threats are logical. There are physical

³ Information and downloads of Design/CPN may be found at <http://www.daimi.au.dk/designCPN/>

threats as well as threats growing out of appropriate policy and system operation. These threats are covered for the most part by the general threats in the Knowledge Base. However, there may be specific issues peculiar to the enterprise under test. These must be considered.

Thus, a complete review of policies and procedures as well as physical security measures must be included. It is reasonable to consider a physical space, such as a computer center, as a security policy domain. When performing such a review, it is a good idea to map the review to an appropriate standard such as ISO 17799 or, in the case of organizations that must comply with specific regulatory requirements and laws, the applicable laws and requirements.

Domain State Change Modeling and Simulation

Once the current state has been modeled and threats against it have been identified, the threats are applied one at a time to the model and the results simulated. The results of the simulation are noted, their impacts analyzed and countermeasures selected and applied. The simulation is then rerun and the results observed.

It is conceivable that the application of a countermeasure or safeguard will have the effect of impeding or stopping the correct operation of the model. In these cases, the selection of the countermeasure must be reconsidered. Equally, the imposition of a countermeasure may have little or no effect upon the protection of the security policy domain against the threat in question. In that case, again, the countermeasure must be reconsidered.

Communications between security policy domains must be modeled carefully to ensure that only permitted communications occur. However, it is not uncommon to find that, although certain communications are undesirable, they are, none the less, occurring. This indicates a discrepancy between the desired, or written, policy and the actual configuration of devices intended to implement the policy.

Final Simulation

Once all countermeasures have been applied and tested at the security policy domain level, the overall enterprise model should be run in the

simulator. It is at this point that individual interactions may cause the model to highlight problems with selected countermeasures or with communications channels not yet considered.

CONCLUSIONS

It is feasible to conduct a risk-based analysis of an enterprise network by modeling the interactions of its security policy domains both with each other and in response to a set of threats against them.

Some of the benefits from applying this type of analysis instead of applying brute-force testing are:

- The enterprise is at less risk of damage from test attacks and penetration attempts;
- Subtle interactions between security policy domains and each other and with threats may show up where they might not under brute force testing;
- It is less expensive to perform modeling than to perform brute-force testing;
- Updating the analysis to accommodate ongoing changes to the enterprise is far easier and less expensive using modeling than it is using brute force testing;
- The ability to manage risks to the enterprise is a direct result of the FARES approach – vulnerability and penetration testing provide only a single component of the risk management process: vulnerability management;
- Since the FARES method is not vulnerability-centric, the need to identify as yet unknown vulnerabilities is not a factor for success;

The risks associated with the FARES method appear to be few. The most important potential risk is that FARES analysts may inadvertently miss an important threat or inter-domain interaction. Other risks relate to the completeness and correctness of the modeling process. Finally, security policy domain identification is critical to the success of the process and the boundaries of these domains may not be clearly defined.

FUTURE WORK

The FARES process has been applied under controlled laboratory conditions but has yet to be used in an uncontrolled environment. Obviously, this is an important next step.

Although the underlying techniques have been proven in actual practice, there may be as yet unidentified subtleties in the FARES process that will not surface until it is applied to a real enterprise. Such field trials are, currently, underway.

As with all modeling processes, there is the opportunity for further refinement. One potential area for investigation concerns the use of other model checkers instead of CPNets and Design/CPN or CPNTools.

Finally, we expect that field trials of the FARES process will suggest further avenues for investigation.

Works Cited

- [AJ02] Jones, Andrew. "Identification of a Method for the Calculation of Threat in an Information Environment."
- [AS91] Anderson, Alison, and Michael Shain. "Risk Management." Information Security Handbook. Ed. William Caelli, Dennis Longley and Michael Shain. 1 ed. New York City: Stockton Press, 1991. 75-127.
- [FS03] Siebenlist, Frank. CORBA-SECURITY-Glossary. 14 Nov 2003 <<http://www-unix.mcs.anl.gov/~franks/CORBA-SECURITY-Glossary.htm>>.
- [FV92] Farmer, Dan, and Weitse Venema. "Improving the Security of Your Site by Breaking into It." 1992
- [PS03] Stephenson, Peter. *Modeling of Post-Incident Root Cause Analysis*. "International Journal of Digital Evidence" Volume 2, Issue 2, <<http://www.ijde.org>>
- [RJ02] Jacobson, Robert V. "Risk Assessment and Risk Management." Computer Security Handbook. Ed. Seymour Bosworth and M. E. Kabay. 4 ed. New York: Wiley, 2002. 47.1-47.16.
- [ST02] Security Tracker Statistics. 2002. SecurityGlobal.net LLC. 23 October 2003 <<http://securitytracker.com/learn/securitytracker-stats-2002.pdf>>.
- [SW00] Sanchez, Luis, et al. "Requirements for the Multidimensional Management and Enforcement (MSME) System."
- [WS98] Security Concepts for Distributed Component Systems. 1. Ed. Walt Smith. 16 June 1998. NIST. 14 Nov 2003 <<http://csrc.nist.gov/nissc/1998/proceedings/tutorB2.pdf>>. (page 53)