

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 data sources, gathering and maintaining the data needed, and completing and reviewing the 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 Service, 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.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) MAR 2010		2. REPORT TYPE Conference Paper Preprint		3. DATES COVERED (From - To) May 2008 – Mar 2010	
4. TITLE AND SUBTITLE IMPACT MODELING AND PREDICTION OF ATTACKS ON CYBER TARGETS. (PREPRINT)				5a. CONTRACT NUMBER FA8750-08-C-0168	
				5b. GRANT NUMBER N/A	
				5c. PROGRAM ELEMENT NUMBER 65502F	
6. AUTHOR(S) Aram Khalili, Brian Michalk, Lee Alford, Chris Henny, and Logan Gilbert				5d. PROJECT NUMBER 0630	
				5e. TASK NUMBER SD	
				5f. WORK UNIT NUMBER 02	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) 21 st Century Technologies 4514 Seton Center Parkway, Suite 320 Austin, TX 78759-5731				8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFRL/RIEF 525 Brooks Road Rome NY 13441-4505				10. SPONSOR/MONITOR'S ACRONYM(S) N/A	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-RI-RS-TP-2010-17	
12. DISTRIBUTION AVAILABILITY STATEMENT <i>Approved for public release; distribution unlimited. PA # 88ABW-2010-0885 Date Cleared: 1-March-2010</i>					
13. SUPPLEMENTARY NOTES This work, resulting in whole or in part from Department of the Air Force contract number FA8750-08-C-0168, has been submitted to SPIE Defense, Security, and Sensing 2010 to be held 5-9 April 2010 in Orlando, FL. If this work is published in proceedings, SPIE may assert copyright. The United States has for itself and others acting on its behalf an unlimited, paid-up, nonexclusive, irrevocable worldwide license to use, modify, reproduce, release, perform, display, or disclose the work by or on behalf of the Government. All other rights are reserved by the copyright owner.					
14. ABSTRACT Modeling an organization's mission vulnerability to cyber attacks requires a description of the IT infrastructure (network model), the organization mission (business model), and how the mission relies on IT resources (correlation model). With this information, proper analysis can show which cyber resources are of tactical importance in a cyber attack, i.e., controlling them enables a large range of cyber attacks. Such analysis also reveals which IT resources contribute to the organization's mission, i.e., lack of control over them gravely affects the mission. These results can then be used to formulate IT security strategies and explore their trade-offs, which leads to better incident response. This paper presents our methodology for encoding IT infrastructure, organization mission and correlations, our analysis framework, as well as initial experimental results and conclusions.					
15. SUBJECT TERMS Cyber Situation Assessment, Cyber Situation Awareness, Cyber Impact Assessment, Cyber Attack Anticipation, Data Fusion, Constraint Satisfaction, Certainty Factors, Cyber Attack Modeling					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 10	19a. NAME OF RESPONSIBLE PERSON George P. Tadda
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) N/A

Impact Modeling and Prediction of Attacks on Cyber Targets

Aram Khalili, Brian Michalk, Lee Alford, Chris Henney & Logan Gilbert

{akhalili|bmichalk|lalford|chenney|lgilbert}@21technologies.com

21st Century Technologies

4515 Seton Center Pkwy Ste. 320

Austin, TX 78759 USA

ABSTRACT

In most organizations, IT (information technology) infrastructure exists to support the organization's mission. The threat of cyber attacks poses risks to this mission. Current network security research focuses on the threat of cyber attacks to the organization's IT infrastructure; however, the risks to the overall mission are rarely analyzed or formalized. This connection of IT infrastructure to the organization's mission is often neglected or carried out ad-hoc. Our work bridges this gap and introduces analyses and formalisms to help organizations understand the mission risks they face from cyber attacks.

Modeling an organization's mission vulnerability to cyber attacks requires a description of the IT infrastructure (network model), the organization mission (business model), and how the mission relies on IT resources (correlation model). With this information, proper analysis can show which cyber resources are of tactical importance in a cyber attack, i.e., controlling them enables a large range of cyber attacks. Such analysis also reveals which IT resources contribute most to the organization's mission, i.e., lack of control over them gravely affects the mission. These results can then be used to formulate IT security strategies and explore their trade-offs, which leads to better incident response. This paper presents our methodology for encoding IT infrastructure, organization mission and correlations, our analysis framework, as well as initial experimental results and conclusions.

Keywords: impact assessment, cyber awareness, mission assurance, information assurance, computer network defense, cyber network operations, simulation, risk management, cyber security, cyber attack

1. INTRODUCTION

Cyber infrastructure and cyber warfare are increasingly important in modern societies. The cybercrime volume continues to increase, yet we only have a basic understanding of the complexities and interconnections of cyber assets. Congress responded to this by passing the Cybersecurity Enhancement Act of 2010,¹ and the federal government's Department of Defense is creating a new cyber command.^{2,3} Our work aims to improve cyber situational awareness and mission assurance by connecting information on cyber assets with their role in an organization's mission. In terms of Endsley's three level situational awareness framework,⁴ our approach provides Level II (Comprehension) and some Level III (Prediction) situational awareness.

Previous work has developed valuable tools and methods that give IT administrators the ability to configure and audit their infrastructure, and probe for vulnerabilities. Many of these efforts, though, work at a very detailed level, at an individual system or service level, and do not expose the implications of their findings. By linking the operation or mission to the cyber resources, our approach provides the ability to assess the mission risk of particular attacks and specific assets' value to completion of the mission.

Consider the operations of a printing business and its use of cyber resources to accomplish its mission. We will use this example throughout the paper. A list of their IT infrastructure vulnerabilities can give the business an idea of how much effort is necessary to address the problems, but it does not give an assessment of the mission risk while the infrastructure vulnerabilities are still in place. For instance, a list of vulnerabilities does not indicate whether their Internet connection is more valuable to their mission than their accounting system.

Our goal is to unveil the dependence of an organization's mission on the IT infrastructure and its resources. We will refer to IT resources throughout the paper. This term includes computers (real and virtual, server and workstation), printers, routers, appliances, smart phones, and other resources. We also use the terms 'cyber' and 'IT' interchangeably. To this end we have developed IMPACT, a tool for Impact Modeling and Prediction of Attacks on Cyber Targets.

2. RELATED WORK

IT security has been extensively studied and many tools have been created to provide or enhance defensive capabilities. We know of no project similar in scope to IMPACT, but there are techniques or tools that are related, have inspired or can mutually benefit from our work.

2.1 Intrusion Detection

Intrusion detection⁵ is the practice of detecting malicious behavior (i.e. violation of security policy) in computer systems or networks. Intrusion detection works at a different level than our work, level I in Endsley's framework. While IMPACT does not work at this level, it can incorporate network status information from an intrusion detection system in its network model, and thus provide current or near real-time awareness with its display and analysis of an evolving situation. A plugin API (Application Programming Interface) is included for this purpose.

2.2 Penetration Testing

Penetration testing⁶ is an authorized attempt to breach security policy in order to discover or verify vulnerabilities in an IT infrastructure. Its goal is usually to identify vulnerabilities so that they can be removed or mitigated. Our approach simulates penetration testing in its network model. However, a simulation is only as good as its assumptions, i.e. the algorithms cannot find any vulnerabilities not in their network model, and hence cannot replace penetration testing as a security tool. On the contrary, penetration testing can be helpful in building a network model, since it aims to find all its vulnerabilities.

2.3 Attack/Protection trees

Attack/Protection trees⁷⁻⁹ are a tool to evaluate different attack or protection methods with regard to a particular target or objective, and differentiate the methods with respect to success chance and resource cost. An attack tree represents an attack plan and all its possible intermediate steps. At the top of the tree is the ultimate goal or root goal. The children of a node are subgoals that enable completion of a parent goal.

The protection tree is the dual of the attack tree. For each attack tree node, there is a corresponding protection tree node that seeks to prevent the attack goal. When a number of subgoals are necessary to enable a particular goal, then the defender need only prevent any one of them. However, if any of a number of subgoals is sufficient to enable the parent goals, then the defender must prevent all of those subgoals simultaneously.

Attack trees relate the various possible ways of achieving a goal, and allow a comparison between them. Components of an attack tree goal include a probability of success, assessment of the impact on the target, and a resource cost. These values must be known a priori for the leaves of an attack trees, though some can be propagated upwards to the other nodes in the tree, e.g. the cost and impact of a node reflect the cost and impact of the node's subgoals. Once all the probabilities have been assigned and cost and impact are propagated up the tree, a cost factor can be calculated, which serves as a basis for comparison. The cost factor is an indicator of the degree of impact achieved for a certain cost at a particular success probability. The path (or multi-path) from a leaf (or leaves) to the root that has the highest cost factor sum indicates the path that provides the most impact for a given cost. Protection trees have cost factors as well, and similar to an attack tree, a protection tree can be used to identify the tactics for a defender to achieve maximal protection value.

IMPACT and protection trees can benefit each other in three ways. Protection trees need to be created from different attack paths that make up the tree. Our algorithms find attack paths through a network and can be used to construct protection trees. Protection trees need an assessment of the impact of a successful attack on an organization. Our work provides these, too. IMPACT seeks to give decision support on allocating defensive priorities and resources, but currently it provides a qualitative result (priority list). Combining our work with protection trees can give a quantitative cost vs. benefit analysis.

3. ARCHITECTURE

Since IMPACT is geared towards levels II and III in Endsley’s situational awareness model, it does not make its own observations. Our framework relies on traditional network data sensors, such as network monitoring or intrusion detection systems. From these a network graph or model is built that describes the IT infrastructure. Currently the models are built by hand, though a plugin API exists to accept updates from sensors. This could be used to start with a minimal model and acquire the rest via updates, if the sensors are sufficiently capable.

To provide a mission risk analysis for cyber attacks, our framework must understand the organization’s mission at an appropriate level of abstraction. To do this, we create uses the business model. We know of no software tools that define or build business models, hence our models have been built by hand. The business model contains the overall mission of the organization, the contributing resources and procedures, and the dependencies between those. Our approach analyzes these to arrive at an overall efficiency score that assess how efficient, based on the maximum efficiency, an organization operates under particular conditions.

For the analysis of mission risk from cyber attacks, the framework must be able to relate the network and business model. This is accomplished with the correlation model. The correlation model associates business resources and procedures with IT resources that support them, such that if an IT resource becomes compromised or unavailable, its effect on the business function can be traced.

These three models: a network model, a business model, and a correlation model form the core of IMPACT.

3.1 Network Model

The network model is similar to network representations found in many monitoring, management and visualization tools. It represents all relevant entities in the network, i.e. IT resources (cf. Section 1). Figure 1 shows an example network for the printing business mentioned in the introduction. The business has a local area network that connects all IT resources except the point-of-sale and the sign printer. Those may be unconnected because they are too old to include LAN connectivity. A router connects that infrastructure to the Internet. Also connected to the Internet is a hosted website, which is under different administrative control (by the hosting company), but still an important part of the business’ IT infrastructure.

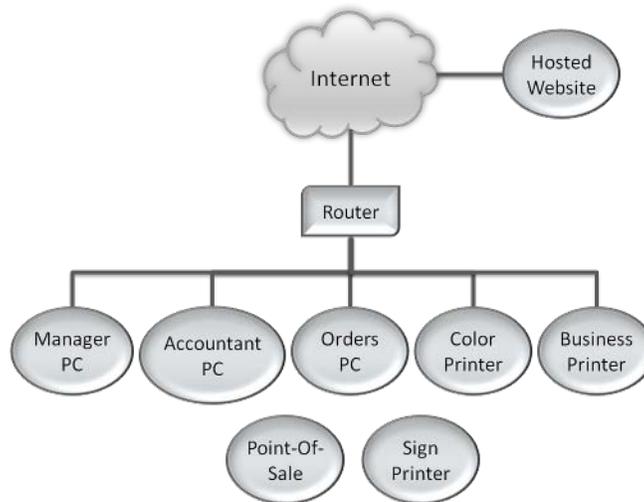


Figure 1. Sample Network Model

Augmenting the physical infrastructure representations are configurations and capabilities. Presently, configurations are the software or firmware loaded or operating on the devices. This includes version number, and may include particular configuration information specific to a software or firmware item on an IT resource. Capabilities are both physical characteristics, e.g. CPU type, memory capacity, etc, and service characteristics. Service

characteristics include the function of the IT resource, e.g. domain controller, file server, development platform. Physical capabilities are currently unused in analysis algorithms, but are intended to match against complex attack requirements in the future. Currently they can be used to display detailed information about IT resources. Service characteristics are included in the network model to accommodate multi-stage attacks that involve local infrastructure.

Consider a phishing attack¹⁰ for internal passwords. An unsuspecting user may be redirected to a fake login screen and prompted to type in their user credentials. If the fake page is hosted outside the company, an application-level firewall may notice that traffic to an internal service should not leave the organization’s network and hence drop the connection. In such a case, an attacker could try to set up a fake login page on a webserver inside the organization’s network, gather credentials, and exfiltrate them later. A webserver is necessary to host the fake login page. Using the service characteristics in the network model, an analysis algorithm can check whether a webserver is present among the resources the attacker has access to and note whether internal phishing is a vulnerability an outside attacker can exploit in a particular scenario.

The network model also includes logical groupings, such as subnets, and vulnerability and exploit information. Vulnerabilities and exploits each have requirements and capabilities. Vulnerability requirements are keyed to particular configuration items, and provide certain capabilities. For example, in a buffer overflow, the capability provided in the vulnerability is the maximum space for the buffer overflow. Exploit requirements need to be matched with a vulnerability’s capability to make it viable to use them together. Exploits also provide capabilities, such as arbitrary code execution, privilege escalation, reading or modifying an arbitrary file, etc.

3.2 Business Model

To assess the effect of cyber attacks on a business or mission, one first needs to have an understanding of the business or mission. To that end, IMPACT requires a business model that describes an organization’s mission. One of the sample business models we built is that of a printing business (see Figure 2). The business makes signs and banners on its sign printer, and glossy brochures on its color printer. Both printers need materials to process, employees to operate them, and orders to be fulfilled in order to generate revenue. These are the major business resources of the printing business.

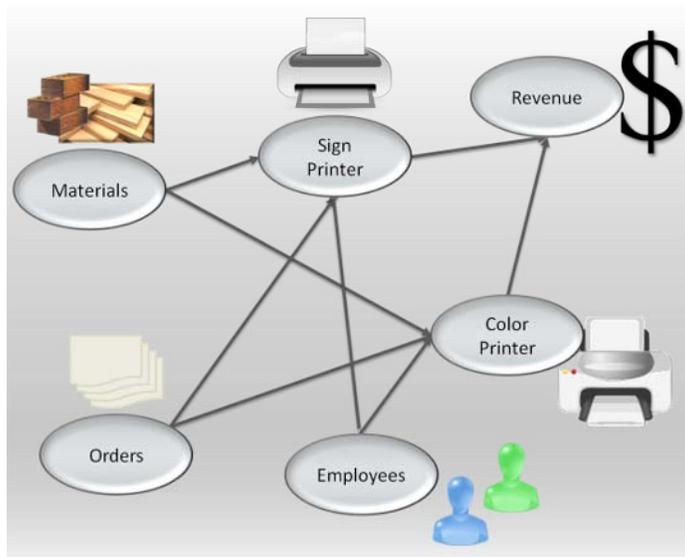


Figure 2. Sample Business Model with causal/temporal Dependencies

Resources in the business model are connected by two kinds of relationships. The first is a causal or temporal dependency, e.g. when ink and paper (or other materials) have to be available at the printing business for anything to be printed. The second relationship is that of an alternate resource or procedure. Organizations often have

Business Procedure	Required Resources	Efficiency
Order Materials	Order online: Management PC, router	100%
	Order online: any PC, router	80%
	Order by phone: <empty>	60%
Online Orders	Order via website: website, Orders PC, router, business printer	100%
	Order via website: website, any PC, router, business printer	80%
	Order via website: website, any PC, router	60%
Store Orders	Order via POS: POS, business printer	100%
	Order via hand: <empty>	60%
Pay employees	Direct deposit: Accounting PC, router	100%
	Print cheques: Accounting PC, business printer	80%
	Write cheques by hand: <empty>	40%

Table 1. Sample Alternate Procedures

policies about how their processes are to be carried out; mindful organizations have backup plans for when the normal policies cannot be carried out. Alternate resources or materials represent these type of backup procedures or resources. In our example printing business, the standard procedure is to order new printing materials from a supplier over the Internet. When the Internet is unreachable, due to IT resource failure or otherwise, the backup procedure is to order the necessary materials by phone. Some sample procedures and alternates for our printing business are given in Table 1. The numbers given are not representative, not based on any data, and were simply chosen for illustrative purposes.

The primary resource or procedure is usually the best known one to accomplish its function in the organization. Alternate resources or procedures are less efficient. In our business model this is represented by an efficiency score. The primary resource or procedure is assumed to have 100% efficiency, alternates are assumed to have less. This assumption is a simplifying assumption, and not strictly necessary. If there is no object with 100% efficiency for a given resource or procedure, then the organization’s total efficiency under best conditions may be less than 100%, which seems contradictory to us. Our model also accepts more than one object among alternate resources or procedures with 100% efficiency; however, we would question whether two separate resources or procedures are entirely equivalent.

Each business model has a root node that identifies the organization’s mission. In case of the printing business we use revenue. This is another simplifying assumption. We do not use profit (presumably a printing business’ actual mission) because our goal is not an optimization tool for an organization’s operations, it is to provide an analysis tool for cyber attack risk to an existing IT infrastructure, and that risk primarily affects revenue rather than profit. This does not mean, however, that IMPACT cannot be used to optimize IT infrastructure to minimize cyber attack risk; in fact, we intend it to be just that.

3.3 Correlation Model

The correlation model is a mapping between a network and a business model, and gives the dependencies of business resources and procedures on IT resources. This section describes the design of the correlation model. Not all features are currently implemented.

The model is used in the business analysis algorithm to check which business resources and procedures are affected by cyber attacks or other IT failures. The model is designed to allow different kinds of relationships to be expressed. A business procedure can depend on zero or more IT resources, and when one resource is

unavailable, the procedure is unavailable. A business procedure that depends on zero IT resources is unaffected by cyber attacks. Similarly, an IT resource can support zero or more business procedures. An IT resource that supports zero business procedures may be a legitimate part of the infrastructure, or it may be a legacy or unused piece of equipment. A side benefit from IMPACT's modeling requirements is the discovery of resources that do not contribute to the organization's mission.

More complex dependencies involve redundancies or graceful degradation. In the case of total redundancies, the business procedure stays available until all redundant IT resources are compromised or have failed. Graceful degradation works similarly, except that the base efficiency of the business procedure is reduced by each successive failure in the dependencies.

Partial redundancy is more difficult to model. When resources become partially unavailable, it is often not clear whether the crucial part is still available. Consider a distributed database where some, but not all, data is replicated across different servers. The model is unable to tell whether a particular query's answer will lie on resources that are still available, or on others that have been compromised or failed. In these types of cases, we suggest using the graceful degradation approach, possibly also in combination with a total redundancy component.

4. IMPACT

The data structures explained in the previous section support our analyses, which provide the insight into the mission risk from cyber attacks. Our framework performs two types of analyses, a network and a business analysis. The network analysis simulates an attack on the IT infrastructure, and identifies its weak and tactically important points. The business analysis starts with a presumed cyber attack or other IT resource failure, and finds the overall business efficiency under that particular scenario. Iterating over all IT resources or even all subsets of IT resources can identify the most crucial resource or set of resources to the organization's mission. In combination, both models can assess the risk of a current cyber attack to the organization's mission, by first finding the vulnerabilities the attacker can exploit in the network model, and then assessing the impact of their unavailability in the business model.

4.1 Network Analysis

IMPACT's network analysis uses the network model and information about vulnerabilities and exploits to assess the network risk stemming from a particular attack. An attacker is assumed to gain initial control of a particular IT resource. From there, the analysis is to check reachable IT resources, gather their configurations and vulnerabilities, and compare them to currently available tools (e.g. exploits or other capabilities), and determine whether all the requirements to exploit a particular vulnerability exist. If so, the identified IT resource is assumed to be compromised and added to the attacker's available resources and capabilities. This procedure is repeated until no additional IT resources can be compromised.

The result is the set of IT resources that can be compromised from a particular starting point, constrained by a particular set of exploits or capabilities. This analysis can be done for every IT resource in the network model, which will establish the set of compromisable resources from each starting resource. Additionally, our tools can make note of the attack paths, i.e. which IT resource led to a compromise of another IT resource. This leads to the discovery of IT resources that are on many different attack paths, and hence act as springboards for many other attacks once they are compromised. IMPACT generates reports that show both metrics.

4.1.1 Example Network Analysis

Consider again the printing business network model from Figure 1. Suppose a customer created an order via the website and attached an executable file that masquerades as a compressed archive, and installs a malware worm program before it opens the included archive. An employee may run the masquerading executable from the Orders PC, where orders from the website are checked. To assess the risk from this attack to the organization's mission, IMPACT would first run a network analysis to determine the extent of IT infrastructure that can be compromised from that particular starting point.

To start this analysis, our framework needs to know the capabilities of the attack, i.e. which resources does it have, and which exploits does it have access to. One initial resource is the Orders PC, where the malware started executing. There may be others outside the printing business (e.g. download server for further malware, but the current network analysis does not support that). The available exploits can range from a particular set to all known exploits, or even an assumed 0-day exploit (see Section 5.1). Suppose the attack was aimed at the print business' router, since they are usually only maintained by specialized staff, which a small printing business may have outsourced, and maintenance may only occur sporadically.

In this case the attack may have access to all known router exploits, perhaps even an unknown 0-day exploit. IMPACT compares these exploits to the IT resources reachable from the Orders PC, and finds a match in the router. The algorithm assumes the attack succeeds, and adds the router to the resources available for the attack. Since no other resources in the network model should be running router software or firmware, the attacks ends there.

4.2 Business Analysis

The business analysis evaluate a business model for its efficiency. It starts at the leaf nodes of the tree rooted at the root node, and selects from the leaf nodes the alternates with the highest available efficiency. Once the leaf nodes' efficiencies are determined they are propagated up the tree to the interior nodes, and finally to the root node. When a node depends on several other nodes, its children's efficiencies together combine to make up that node's efficiency score. We are unsure of the best way of doing this, so our framework offers a number of options in aggregating child node efficiencies:

- **low watermark** the lowest of the efficiencies
- **arithmetic mean** the sum of the efficiencies divided by their number
- **product** the product of the efficiencies,
- **geometric mean** the n th root of the product, where n is the number of children
- **weighted average** the sum of the efficiencies after multiplication by weights that add to 1.

To evaluate the business model during a cyber attacks, another step is added. Whenever a resource or procedure is selected in the business model, its requirements in the IT infrastructure are checked for availability. If all requirements are available, the node is used in the analysis; if not, an alternate must be chosen. If no further alternate exists, a default node with an efficiency of 0% is applied.

4.2.1 Example Business Analysis

Consider again the cyber attack from Section 4.1.1. The Orders PC and the router are compromised, and none of them have redundancies or degrade gracefully. Consult Table 1 for available business procedures. The online materials order are unavailable (as the router could intercept account number information or redirect payment authorization somewhere else), so ordering becomes a 60% efficient process. Retrieving online orders becomes completely unavailable, and has 0% efficiency. Store order remain unaffected at 100%, while employee payment goes to check printing at 80% efficiency. According to the business model in Figure 2, both the Color and Sign Printer depend on Materials, orders (online and store) and employees. If we choose the arithmetic mean type of efficiency aggregation, each of the printers can operate at $\frac{60+0+100+80}{4} = 60\%$. The print business revenue depends on both printers. Suppose the aggregation at the root node is multiplication. In that case printing business revenue is still $0.6 \times 0.6 = 0.36$ or 36% efficiency. In this case the efficiency cost of the router and the Orders PC is 64% (100% - 36%). The alert reader will notice that the router usage dominates the Orders PC usage in the business procedures, and hence does not contribute separately to the efficiency cost, i.e. the efficiency of a compromised router alone is also 64%. Currently, our framework does not especially identify efficiency cost by subgroups of compromised resources, or reduce the set to the smallest equivalent, but proper (i.e. repeated) usage of IMPACT will nonetheless expose these types of relationships, even when they are not obvious in the models.

5. CYBER AWARENESS

As a result from the analyses described in the last section, an organization can gain *cyber awareness*. The network analysis shows which IT resources are of tactical advantage in a cyber attack. The business analysis shows how a particular cyber attack or other IT failure can affect an organization's ability to accomplish its mission.

These analyses can be applied as a planning tool ahead of any incidents, or in situ as a decision aid. As a planning tool our approach allows evaluation of current and potential IT infrastructure in terms of the potential for attack that they offer possible attackers (network analysis) and in terms of risk posed to the organization's mission (business analysis). The first part helps network administrators prioritize their attention, as it identifies IT resources that can be of high utility to possible attackers. A number of reasons can cause an IT resource to be identified as high network risk, and it can be dealt with in different ways, e.g. by making it a high priority to keep the IT resource current with all patches and configuration recommendation, by changing its location in the network topology, or perhaps by splitting the functions among several different IT resources.

The second analysis helps business analysts and network architects find weaknesses in their IT strategy, as it identifies IT resources whose absence or lack of trustworthiness can cause large disruption to the organization's mission. After a resource has been identified as high mission risk, a network architect or business analyst can take steps to mitigate the situation, e.g. by building redundant services that stay available even when some parts are compromised, or by devising alternate means (and IT resources) by which the same mission goals can be accomplished.

As a decision aid during an incident, our framework can analyze the options a current attacker has (based on best current knowledge) and project their network and mission risk. Additionally, IMPACT can analyze and project the impact of defensive measures on both the attackers options inside the network and the effects on the organization's mission ability.

5.1 0-day vulnerabilities

In our current state of software practices, new vulnerabilities continue to be found, and while they are sometimes brought to the attention of the organization maintaining the software or IT resources (appliance, router, printer, etc), often they are not and instead are used to craft new attacks that are very difficult to plan for using current IT security tools. This issue has generated significant interest in capabilities to deal with 0-day vulnerabilities. Our approach does not detect vulnerabilities or attacks, but it has the ability to assess the potential impact of a *particular* 0-day vulnerability.

The network model includes a set of vulnerabilities and exploits that are based on current knowledge and need to be maintained to be kept up to date. However, it is possible to add a vulnerability that is not known to exist and re-run various analyses on the resulting configuration. If, for example, a network administrator wants to assess the risk of a new vulnerability in SQL Server or BIND (Berkeley Internet Name Domain), she could create such a vulnerability in the network model, and all further analyses would assume that such a vulnerability exists and project the impact onto the network and business models.

Similarly, if analysis of a current attack leads to the discovery of a previously unknown vulnerability, the network model can be adapted and analyses re-run to assess the risk posed by the new vulnerability.

6. FUTURE WORK

IMPACT is a work in progress, and much remains to be done. Central to our approach are the network, mission and correlation models, and each of them can be made more expressive. In the network model, more decision aid functionality can be added, e.g. the automatic calculation of the network cut that leaves the largest portion of the IT infrastructure intact and separated from the attacker (as is best known).

One of our goals is to model complex attacks. Our current attacks are single stage attacks, while real-world attacks have become very sophisticated and multi-pronged. Some support for more complex attacks exists (cf. resource types in the network model), but critical parts are not yet implemented.

The description of the correlation model (cf. Section 3.3) already mentions difficulties expressing redundancy. The correlation model best deals with a 1:1 mapping of IT resources to mission resources and procedures. Other mappings are possible, but some many-to-one (redundancy) and many-to-many relations are not properly dealt with.

Our long term goal is move IMPACT beyond the point of a decision aid and into active response. The next step in the evolution of computer systems after aiding a human operator in their decision is for the system to make limited decisions itself. IT security can be a split second challenge, and reliable automated response systems can bolster the defense of IT resources.

7. CONCLUSIONS

IT infrastructure risk has been extensively studied, though it is far from a solved problem. To our knowledge, this is the first work tying this risk to the underlying reason for the existence of the IT infrastructure – namely the organization’s mission. We created an approach that allows for the definition of the organization’s mission, its dependence on IT infrastructure, and the analysis of the mission risk that comes from this dependence. It can be used as a planning tool to evaluate and optimize an IT infrastructure, and as a decision aid to evaluate potential action under an attack.

8. ACKNOWLEDGMENTS

This research and development effort is supported by the Office of the Secretary of Defense (OSD) and the Air Force Research Laboratory (AFRL) through a Small Business Innovative Research (SBIR) grant, under contract number FA8750-08-C-0168. We thank program manager George Tadda for his encouragement.

REFERENCES

- [1] “The Cybersecurity Enhancement Act of 2010.” H.R. 4061, 111th Congress.
- [2] Miles, D., “Gates Establishes New Cyber Subcommand.” American Forces Press Service (June 2009). <http://www.defense.gov/news/newsarticle.aspx?id=54890>.
- [3] “24th air force activated, 2 units realign in joint ceremony.” AFNS (August 2009). <http://www.af.mil/news/story.asp?id=123163831>.
- [4] Endsley, M. R., “Toward a theory of situation awareness in dynamic systems,” in [*Human Factors*], **37**, 32–64 (March 1995).
- [5] Denning, D. E., “An intrusion-detection model,” *IEEE Transactions on Software Engineering* **13**, 222–232 (February 1987).
- [6] Scarfone, K., Souppaya, M., Cody, A., and Orebaugh, A., [*Technical Guide to Information Security Testing and Assessment*], ch. 5, Special Publication 800-115, NIST (September 2008).
- [7] Weiss, J. D., “A system security engineering process,” in [*Proceedings of the 14th National Computer Security Conference*], 572–581 (1991).
- [8] Schneier, B., “Modeling security threats,” *Dr. Dobb’s Journal* **24** (December 1999).
- [9] Edge, K., Raines, R., Grimaila, M., Baldwin, R., Bennington, R., and Reuter, C., “The use of attack and protection trees to analyze security for an online banking system,” in [*Proceedings of the 40th Hawaii International Conference on System Sciences*], 144b (January 2007).
- [10] Jacobsson, M. and Myers, S., [*Phishing and Counter-Measures*], John Wiley and Sons, Inc. (December 2006).