

Evaluating the Impact of Cyber Attacks on Missions

Scott Musman, Aaron Temin, Mike Tanner, Dick Fox, Brian Pridemore

MITRE Corp, McLean, VA, 22102

Abstract

Using current methods, it is virtually impossible to determine the impact of a cyber attack on the attainment of mission objectives. Do we know which mission elements are affected? Can we continue to operate and fulfill the mission? Should we wait for recovery? Can we salvage part of the mission? Since it is currently so difficult for humans to comprehend the mission impact of a cyber incident, our ability to respond is much less effective than it could be. We believe that improved knowledge of the mission impact of a cyber attack will lead to improved, more targeted responses, creating more attack resistant systems that can operate through cyber attacks.

Our work addresses the “mission” part of “mission assurance,” focusing on cyber mission impact assessment (CMIA). Our challenge is to create mission models that can link information technology (IT) capabilities to an organization’s business processes associated with Measures of Effectiveness and Performance (e.g., attrition of enemy forces, targets destroyed, blue force protection). Measuring mission impact requires knowing the mission activities that fulfill mission needs, the supporting cyber assets, and understanding how the effects of an attack change mission capability. This paper is about developing the techniques that make estimating the mission impact of cyber attacks possible.

1 Introduction

Increased integration of computers and information technology (IT) into the war fighting process has created an environment where compromise, damage, or loss of IT assets can result in mission failure. Thus, our expectations of the success of a mission that is under cyber attack (i.e., attacks against the IT supporting a mission) greatly depend on the understanding of how the cyber attack has degraded capabilities in kinetic (non-IT) space. This report on research in progress describes a system that evaluates the effect of a cyber attack by predicting the impact of the attack on the mission’s measures of effectiveness.

Consider the following example. A time sensitive targeting (TST) mission thread is being executed. The mission commander’s ability to select a weapon to deploy against a target depends on information about available airborne weapons and ground-based weapons. Either type of weapon might be useful depending on the type of target. At 1130 hours, two events occur:

- The TST cell is unable to access airborne weapon data for unknown reasons.

- The network operations center supporting the TST cell has a report of a denial of service (DoS) attack on several routers, disrupting network traffic. There is no estimate of how long the effect of the attack will last.

Currently, the mission commander only knows that there is no access to some of the data needed. The best he can do is to proceed with the data to which he has access, which means planning an attack using a ground-based weapon, even if an airborne weapon would improve the mission's chances of success. At the same time, the IT administrators know there is a cyber attack that is stopping data from moving over the network, but don't really know which mission systems and missions rely on the routers under attack. The IT administrators are unable to tell the mission commander what he needs to know to figure out his options.

We propose that information about the DoS attack be provided as input to a model of the mission and its supporting IT assets. The model infers that the routers under attack are necessary for transmitting the airborne asset status to the TST cell, confirming the connection between the two events. Further, the model calculates the change in the measures of mission success as the duration of the attack extends, and provides the mission commander with the assessment that he can choose to wait 10 minutes for the attack to be cleared without reducing the likelihood of mission success; and if at that point the attack can be cleared within another 10 minutes, he should wait for connectivity to be restored, otherwise he should continue the mission with just the ground-based weapons information.

2 Previous Work

We currently have very little capability to estimate the mission impact of cyber incidents. The state of the art is such that even the most basic mapping of dependencies among mission objectives, mission activities, and IT assets rarely exists. Even when these mappings are determined as part of risk analysis, they are not carried over into use operationally when incidents occur.

The thesis of Fortson [2007] highlights a number of deficiencies of current practice, describes a number of scenarios illustrating operational deficiencies, and provides requirements for an impact assessment solution. Although not phrased in the following manner, the various examples highlight the objectives for mission impact assessment:

1. Make it possible to document the dependency relationships between cyber assets, mission activities, and mission objectives so that the relationships can be used operationally.
2. Make it possible to determine the mission impact of the cyber attack based on the timing and duration of the incident.
3. Make it possible to predict mission impact, even if an affected IT resource is not currently in use.
4. Make it possible to predict the impact on mission instances that are planned or anticipated in the future.

These objectives impart requirements and restrictions on appropriate techniques for a solution. For example, understanding the relationship between incident duration and impact requires that time versus mission value knowledge about activities and data be captured in mission models. These temporal mission system characteristics are rarely considered or documented. Popular architectural notations such as UML, or DoDAF do not include these aspects in their descriptions of the mission.

Existing practices in modeling mission systems are inadequate for our intended purpose. Some existing modeling approaches (e.g., UML, DODAF) are diagrammatic rather than computable. Some modeling paradigms lack the ability to represent time, or workflows (e.g., Bayesian networks, Influence diagrams, dependency maps), which is important since the duration of an incident will often affect the amount of impact an incident will have. Critically, many approaches are also unable to represent information dependencies, a necessity since information attacks are common in the cyber domain.

Although there are existing tools for performing cyber risk assessments [Watters, Web][Whiteman, 2008], the mission models that are created and used in these tools have limited use for computing online impact assessments. Because formal cyber risk assessments are currently an offline process, they focus on the potential cyber effects against a wide variety of possible mission instances, where the specific timing and duration of the attack effect is not specified. As a result, risk assessment mission models tend to lack, for example, timing and workflow information which make it impossible for them to differentiate between attacks that can be recovered from quickly and attacks that would take much longer to recover from. Since many missions are dynamic and have temporal constraints the added details that we would include in our mission models (e.g. activity timing, workflow, and MOE mappings) will allow us to make additional mission impact assessments that are not possible with a more static model.

Existing DoD processes label systems and information as mission critical that are then considered to be mission critical for all time. The reality is that systems and information can be more or less important depending on the time, or phase of a mission. Knowing which systems and information are important to the mission at the time of an incident can help to focus recovery activity and speed up the recovery process. Atemporal representations of missions cannot accommodate these requirements.

Some related work has focused on battle damage assessment (BDA). Although BDA is an integral part of the military process, as of yet there is no standardized BDA process for IT effects [Thiem 2005]. There is an important distinction between BDA and mission impact assessment (MIA). BDA assesses the damage associated with an attack on the resources, while MIA assesses the anticipated effectiveness of the mission resources to carry out the mission after an attack has occurred. In this context we are equating MIA with combat assessment (as defined in JP-102, 2006). Gramaila and Fortson [2007] focus mainly on the BDA processes, rather than the impact assessment process.

Since an important characteristic of impact assessment is its temporal nature, more than just knowing the current capabilities of the IT and mission activities immediately following an attack, it is sometimes necessary to be able to make predictions; therefore, effective impact assessment

involves understanding anticipated activity. Whether it is scheduled, or based on historical expectations, anticipated activity indicates what is likely to be lost in the face of failure or degradation. It is also worth considering that, depending on where and how your mission systems are instrumented, you may only be able to observe indirect effects of an attack, and so in some circumstances the observed lack of expected activity may be the first opportunity to notice that there is a problem that threatens the mission objectives.

There is a large body of existing work, tools and techniques that address mission modeling [Clancy et al, 1998]. Emerging standards such as Business Process Modeling Notation (BPMN) [White and Miers, 2008], and the integration of the modeling notation with executable simulation engines [Anupindi 2005], provide methods to describe a business process as an executable model that can then be used predict various mission specific metrics and measure various performance characteristics given architectural decisions.

For CMIA we consider the possible effects of cyber attacks, as would be reported by a BDA process. Existing cyber attack effect models such as those discussed by [Howard 1998] are less suitable for our needs. Howard's incident taxonomy is information centric, and hence does not include the process oriented characteristics needed to compute the impact of activity interruption, degradation, fabrication, or information unavailability. Although not focused on mission relevance, Blyth and Kovacich [2006] describe effects that come closest to addressing our needs.

3 Technical Approach

Our technical approach involves dividing the problem into several parts. First, we discuss the requirements for modeling missions. Next, we show how we have reduced the number of cyber attack scenarios for us to consider for making impact calculations down to only six classes describing the effects on IT of any cyber attack. We then discuss how we express knowledge of the mission activities and the supporting IT in business process modeling notation (BPMN) and use that to compute measures of effectiveness for a mission instance.

3.1 Requirements for Modeling Missions

To understand the characteristics of mission systems, we reviewed a variety of mission systems—including advanced sensor networks; time critical targeting; JSTARs; the FAA's enroute automation system; a census address canvassing system; and several others. The result of this review was a list of modeling requirements for how the expressiveness of the mission model affects our ability to compute mission impact (see Figure 1).

- Dependencies between mission elements
 - Allows us to relate between Mission Objectives, activities, cyber assets, information assets
- Workflows
 - Makes it possible to represent ordered interdependencies, and forecast the impact of resources not currently in use
- Uncertainty
 - Allows us to represent the relative likelihood of events, and outcomes
- Utility
 - Represents the value estimates of different mission outcomes, since they may not all be equal
- Time Value Characteristics of activities and information
 - Makes it possible to represent time constraints for activities and information, and predict how the duration of an incident changes its impact
- Fallback and failover activities
 - Represents what kicks in, in the face of failures
- Implicit mission decisions
 - Allows us to capture when certain mission outcomes depend on “built-in” decisions that can change when information is no longer available
- Mission MOE’s/MOP’s
 - We can’t evaluate what we can’t measure
- Scenario characteristics
 - Sometimes the impact of an incident depends on the context of how/where the system is being used

Figure 1: Modeling Requirements

In order to understand how to compute mission impact it is necessary to understand what impact is and how it might be reported. Since impact assessment is concerned with changes in the expected outcome of a mission, the types of impacts can vary. Consider the following examples that illustrate a variety of mission impact assessments:

- We can no longer determine which ground weapons are available
- The mission system is now operating at only 70% capacity
- We can only hit 50% of the high value targets
- Target #356 can no longer be engaged
- Our planned engagement of target #4557 is unaffected by the cyber attack
- There will not be any impact to the mission if we can recover from the cyber attack within the next 15 minutes
- Because we cannot restore the server till tomorrow, tonight’s planned mission will be affected

These impact statements represent different types of impacts on a mission system but are not mutually exclusive. In general, impact statements might be in terms of:

- a) ability to perform mission activities
- b) capabilities of the mission system in general
- c) achieving specific mission objectives
- d) information about specific mission instances
- e) prediction of how mission impact might vary over time
- f) prediction of how affected resources not currently in use may cause future impact
- g) prediction of how affected resources may cause impact on future mission instances

The fact that a mission impact assessment might involve any of these statements illustrates the complexity of the problem of selecting a technical approach to compute it. It would seem necessary to either select a general purpose approach that can compute these different statements, or to pick a subset of these impact statement types that is “good enough” to provide operational value. In either case, we still want only a single mission model over which to compute these impacts. Thus the mission model must be descriptive enough to support the different calculations, whether they are temporal, probabilistic, or value based.

3.2 Representing cyber attacks

Although various languages exist that can be used to describe cyber incidents (e.g., IDMEF, Common Event Expression (CEE), etc.), these languages characterize incidents in terms of the activities of the attack (e.g., a rootkit attack modifies a system library file; or a buffer overflow allows a user privilege escalation). What is needed for impact assessment is a standardized way to characterize the *effects* of cyber incidents, and as of yet no such language exists. As a step to address this deficiency we developed categorical descriptions of cyber attack effects (Figure 2). Interruption, interception, modification, and fabrication were in Blyth and Kovacich; the others we added based on our examination of the mission assurance domain.

- Degradation
 - An attacker causes a degradation in the performance of an information asset
- Interruption
 - An attacker causes an information asset of the system to become unusable, unavailable, lost for some period of time
- Modification
 - An attacker causes a modification of information, data, protocol, or software
- Fabrication
 - An attacker causes information to be inserted into the system.
- Unauthorized Use
 - An attacker uses the system resources for their own purposes.
- Interception
 - An attacker causes or takes advantage of information leaked from the system

Figure 2: Cyber Attack Effects

Regardless of the mechanism used in a cyber attack, we believe that its effect can be characterized as being one or more of the listed effects on one or more of the mission IT assets (including both IT components and information assets). Although our characterization of attack effects is not a formal language, when combined with information about which resources are affected and estimates of start and end times for the incident, these categorizations provide enough information to allow us to compute impact, and provide the basis for the rest of the work we have done on impact estimation.

3.3 Modeling Missions for CMIA

A challenge of cyber mission impact assessment is relating the needs of the mission to the IT that support it, and capturing these relationships in a mission model. We view a mission as being a collection of activities that must be performed to accomplish a task. The activities are accomplished using mission resources, some of which are human, some might be mechanical, and some are IT. Associated with these activities, various types of information are exchanged, created, or used. The specific information involved is typically described as an information asset, where information assets are assigned to resources, some of which are IT.

Based on analysis of model representation versus impact computation tradeoffs we selected BPMN, along with some proposed extensions to represent information dependencies, as the formalism in which to represent mission systems. BPMN is an emerging standard for process engineering, so significant modeling expertise is available. There are COTS products that connect BPMN models to executable simulation engines that support offline performance analysis that is similar enough to some of the impact calculations we need to perform.

Figure 3 illustrates a top-level mission model that, when used in a simulation (e.g., iGraphx software), allows us to modify various mission system characteristics (e.g., add or remove sensors, perform activities faster, etc.) and compute the effect of these changes on simulated mission outcomes (MOE's).




Figure 3: A top-level TST (like) mission model in iGraphx that can compute mission MOE's

This model represents a performance engineering model, in this case for the domain of Time Sensitive Targeting (TST) that may already exist for the domain of the mission system for which we are trying to perform CMIA. Typically missing in existing models that represent mission systems are the IT resource dependencies and the information dependencies.




Figure 4: Information sources for representing IT include network diagrams, resource inventories, plans/schedules, architectural documentation, and subject matter experts

As illustrated in figure 4, the information used to populate the IT portion of the model is derived from several sources. A network diagram usually exists that inventories the network hardware and describes how it is interconnected. Information from network audits and vulnerability testing can also populate this information, as well as identify the software applications that run on the hardware. At the software level, we want to identify algorithmic or configuration items that affect the workflow, such as data caching, or flows decision made by the software algorithms.

Figure 5 illustrates an example of the details found in the sub-model for the “identify engagement options” activity of figure 3. This model fragment shows how the mission workflow has been expanded to include activities for the IT resources. Although not illustrated in the diagram, modeled resources mapped to IT assets are associated with each activity. Also not shown are various modeling parameters associated with activities that identify the timing and decision points that might affect the mission MOE calculations. By explicitly adding the IT related activities into the model, any changes to the capabilities of the IT resources to reflect the effects of a cyber attack can now change the outcomes of running the model. This particular model includes parallel activities to access information about air and ground weapons assets, and also includes a representation of software caching that makes access to remote information unnecessary if recent status information is available locally. Backup and failover activities can also be represented.




Figure 5: Part of an IT activity model for assessing the available weapons assets

The ability to represent information asset dependencies is missing in almost all mission model representations. Existing workflow modeling approaches, such as BPMN, do not include acceptable methods to represent information to activity dependencies in computational form. Such dependencies are typically represented only graphically. Since cyber attacks include data attacks that affect integrity or confidentiality, we want to represent how each modeled activity depends on these information assets. Moreover, we want to do it in a way that allows us to compute the resulting impact on MOE's when the integrity or confidentiality constraints of information assets are violated. These information dependencies are shown in figure 6.

To address this modeling deficiency we have proposed some extensions to the BPMN models to allow us to explicitly represent the constraints in a computable form. By defining information assets as resources, activities that use these information asset resources can be defined in the model. Since it is possible to define attributes (variables) for resources, we will define attributes for integrity and confidentiality for all of our information asset resources. If activities are then performed using information assets whose integrity or confidentiality constraints have been violated, we can have the model implement mathematical functions that would alter the MOE calculations.

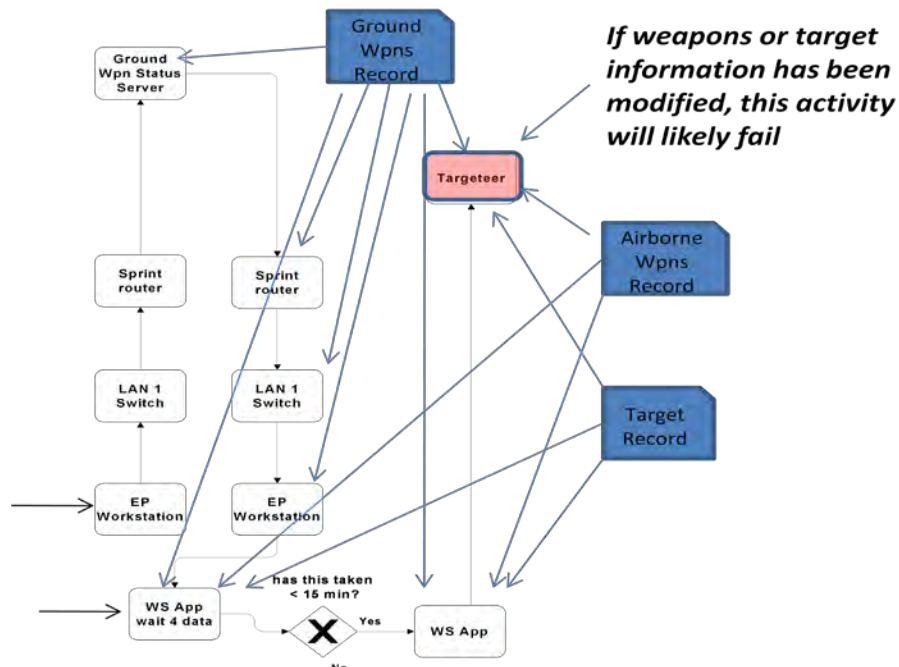


Figure 6: information dependencies associated the IT activities

3.4 Computing Cyber Mission Impact

To estimate mission capability we apply computational algorithms to a mission model in a manner that allows us to link system capability, to mission-oriented MOEs. Using our approach, when a cyber attack occurs, we envisage that an incident report would provide details of the effect on IT resources. We use that information to modify the mission model to reflect changes in system capabilities caused by the cyber attack. Then we rerun the model to produce new MOE estimates. This is illustrated in figure 7. The impact of the attack can then be determined by comparing the two MOE values (before and as a result of the incident).




Figure 7: Estimating mission impact by comparing model MOE's with and without the effects of the cyber attack

Our method currently requires manual intervention to alter the mission model to reflect the cyber effect of the incident, and repeated runs of the simulation to reflect the normal variations in mission instances. Below we illustrate an example of computing mission impact for a specific incident.

Since our objective is to implement a solution that can be run on-line and produce mission impact estimates as incidents occur, our current manual intervention approach is an unacceptable solution. Additionally, simulation as a computational strategy to compute impact is not an ideal solution for producing timely impact estimates. Fortunately, we are aware of several alternatives, and are exploring techniques to transform (or compile) our mission model into constraint models, decision trees, or schedule representations that would be amenable to much more efficient run-time evaluation.

Since there are currently no standardized processes for characterizing and reporting cyber incidents to a system such as the one we are developing, we are also working on how to use our approach offline to evaluate the cyber mission assurance properties of mission systems. Traditional approaches to cyber risk analysis are capable of identifying high risk components, but say very little about which specific threats will have the most impact, how the timing of attacks affects impact, or what to do when an attack occurs. Our techniques might be used to pre-compute a playbook of response options and actions in the face of different cyber attacks.

A cyber incident forces mission personnel into decisions that relate to the executing mission. There are a finite number of alternatives that we need to compute to be able to assist mission operators in making a decision:

- Whether to continue with the mission using the affected IT resources (if possible)
- Whether to continue with the mission and not use the affected IT resources (if possible)
- Whether to continue with the mission after having recovered the affected resources
- Whether to abandon the mission

4 An Example of Computing Mission Impact

Consider the following incident:

11:30 hours: An inability to access AWACS or AOCC (both of whom can provide information about available airborne weapons) coincides with a network DoS reported on several GIG routers. There is currently no way to know how long the DoS attack will last.




Figure 8: Access to remote mission information becomes unavailable due to an attack

This incident is one where access to a source of weapons information is made unavailable, as shown in Figure 8. Since the sources of weapons are independent of each other for any given target, there is some likelihood that there may be ground weapons available to engage the target, and there is a different likelihood that airborne weapons might be available. Some weapons are more suited than others to the particular type of target that is being engaged. Our mission model represents our characterization of how these various factors affect mission outcome.

Since the mission thread can proceed using only the information about available ground weapons, the mission level decision posed to mission personnel is to try to understand: (1) whether they are going to be better off continuing with the mission using only the partial weapons availability information, (2) whether to hold off in proceeding with the mission until after the missing information about available weapons becomes available, or (3) whether to abandon the mission instance because of the incident.

Our mission model makes estimating the impact of these options possible. Figure 9 illustrates estimating a mission MOE for the various decision cases.




Figure 9: Impact curves and conclusion for losing access to airborne weapons status

The green line in the figure illustrates the model estimated impact variation over time when waiting for recovery from the incident to access the airborne weapons status data. Since there would be a delay in proceeding with the mission till recovery takes place, the curve computed by the model reflects the fact that less time to prosecute a time-sensitive target leads to less chance of mission success. The red circle in the figure on the vertical axis illustrates the reduced chance of mission success if the mission was to continue immediately, using only the ground weapons known to be available. The red line in the figure illustrates the temporal characteristics of holding off acting on the partial weapons information. Although displaying results in this form is not ultimately the way we want to present this information to an end-user, these curves can be the basis of advice to mission personnel. These curves indicate to us (as technical professionals) that in this situation the mission personnel can afford to wait 10 minutes to see if the denial of service ends, and suffer no additional impact if they then proceed with the information they currently have now. It also indicates that if they believe the incident can be recovered within 20 minutes, obtaining additional weapons options, then the likelihood of mission success would increase over the options they have right now.

This example of estimating mission impact illustrates the computations we can now make with our mission model when a cyber event occurs, and illustrates putting the results of those calculations in a mission context. The result helps to inform mission personnel in making operational decisions about what to do as a result of the cyber attack.

5 Summary

We have demonstrated how one can compute the mission impact of cyber attacks and shown how the outcomes of the impact estimates provide mission personnel information relevant to mission level decisions. The models shown in Figures 3 and 5 produce the graph in Figure 9, which is the complete impact of an example incident. We have not addressed the question of how to present this information to mission operators, although we are fairly sure that the graph shown is not the presentation we would like to use.

Our work on characterizing cyber attack effects is pertinent to all related activities that need to report cyber incidents. The DIMFUI effects (Figure 2) at least notionally provides a set of output statements for a cyber BDA process. With attacks characterized in this way, a number of mission-level assessments can be supported, one of which is the mission impact described in this paper.

Our analysis for understanding what to include in mission models provides a set of requirements for the documentation and characterizing of mission systems in order to be able to do mission impact assessment (Figure 1). It is important to note that these things are necessary for assessing mission impact of cyber events, or any mission resource setbacks, irrespective of if there is a mission model to support automated assessment.

This paper describes research in progress, and our work continues on developing dedicated software that can do the calculations that we have so far demonstrated by making manual changes to the models.

6 References

Air Force Operational Test and Evaluation Center (AFOTEC/XRC). (1995). AFOTEC 99-101, *Test Concept Handbook*, Jan 95. Kirtland AFB, NM.

USAF (USAF/TEP). (1994). AFI 99-103, *Test and evaluation process*, 25 Jul 94. Washington, DC

Anupindi, R. "Managing Business Process Flows: Principles of Operations Management", 2005, Prentice Hall, ISBN [0131676865](#)

Blyth A., and Kovacich G., "*Information Assurance (security in the information environment)*", 2nd edition, 2006. Springer-Verlag Ltd, London

Clancey, W. J., Sachs, P., et al. . Brahms: "*Simulating practice for work systems design.*", Int. J. Human-Computer Studies, 49: 831-865, 1998

DeMarco, T. "Controlling Software Projects: Management, Measurement, and Estimates," Prentice Hall, 1982 (ISBN 0-13-171711-1)

Grimalia M, and Fortson L. “*Towards an Information Asset-Based Defensive Cyber Damage Assessment Process*”, Proceedings IEEE Symposium on Computational Intelligence in Security and Defense Applications (CISDA 2007), 2007

Howard, J., Longstaff, T. . “*A Common Language For Computer Security Incidents*,” Sandia National Laboratories, Sandia Report, SAND98-8667 1998

Fortson L. “Towards the Development of a Defensive Cyber Damage and Mission Impact Methodology “, AFIT Masters Thesis, March 2007

Lane, N.E. (1986). Issues in performance measurement for military aviation with application to air combat maneuvering (NTSC TR-86-008). Orlando, FL: Essex Corporation.

Theim ,L. “A Study to Determine Damage Assessment Methods or Models on Air Force networks,” Department of Engineering and Management, Air Force Institute of Technology, Wright Patterson Air Force Base, OH, 2005.

TST, “*MULTI-SERVICE TACTICS, TECHNIQUES, AND PROCEDURES FOR TARGETING TIME-SENSITIVE TARGETS*”, 2004, FM 3-60.1, MCRP 3-16D, NTTP 3-60.1, AFTTP(I) 3-2.3

Watters, J., “*RiskMAP — Tool for building a business case for investing in security*”, Web, <http://www.thei3p.org/publications/>

White, S., and Miers, D. “*BPMN Modeling and Reference Guide*”. Future Strategies Inc.. 2008, [ISBN 978-0-9777-5272-0](https://www.futuresstrategies.com/publications/BPMN-Modeling-and-Reference-Guide.pdf).

Whiteman, B. 2008, “*Network Risk Assessment Tool (NRAT)*”, IA newsletter, Vol 1, Spring 2008, http://iac.dtic.mil/iatac/download/Vol11_No1.pdf