Information Technology Security Report, Vol 4, No 3 (1999) 54-62

Achieving Information Resiliency

Paul Zavidniak, Dr. Anita D'Amico, and Dennis H. McCallam

Introduction

Information resiliency refers to the continuous availability of uncorrupted mission-critical information to support business or military operations, even under the threat of a cyber attack. An information resilient enterprise will continue to engage in its critical operations, despite the attacker's attempts to intrude, corrupt or deny service. The manner and efficiency with which the operations are conducted may change somewhat, but they remain operative. The commercial world needs information resiliency to maintain its computing operations in order to prevent financial losses, while the military needs it to prevent casualties and tactical losses. In addition, information resiliency is needed to ensure that a country's critical infrastructure (e.g. transportation, financial industry, electrical power) continues to operate during hostile attacks against their computing and communications systems.

To achieve information resiliency during an Information Warfare (IW) attack, one must understand how observed security breaches (e.g. password failures, intrusion detection, unusual network activity) fit into a bigger picture. By knowing whether a single security event is just a spurious action, or is part of a larger IW campaign, the Information Protection Manager can more appropriately take actions to thwart the attack, respond with countermeasures, and prepare to recover the information and communication systems.

To understand the big picture, and ultimately to achieve information resiliency, one needs an overall model of the cyber attack, from the initial probing by an attacker, to the actual attack launch, to how the target system responds. The model must consider the attacker's (offensive) actions and the defender's actions, as well as the mission impact. In this paper, we present one such model, referred to as the IW Timeline.

The objectives of this paper are to:

- present an IW Timeline model of a cyber attack cycle, and
- offer a strategy for changing the timeline such that defensive tactics used against attacks can become more proactive and less reactive.

Since our timeline model and our initial analyses were born out of military IW, we will start with a brief review of concepts we borrowed from that area.

Reaping the Benefits of Military Approaches to IW

There are many similarities and a few differences between military IW and commercial IW. The US military's pursuit of IW concepts has been somewhat more organized and broader in scope, due to the high priority given to Information Superiority by all of the services, the intelligence agencies, and the Joint Chiefs of Staff. We have found several aspects of military IW to be particularly useful to us in framing our IW Timeline model. We briefly review them below.

The Warfare Mentality

First, the military's warfare mentality facilitates the framing of information

security policy and tactics whether one is in the military or not. While the US military may have curtailed its use of the term "Information Warfare" in preference for the more politically palatable terms "Information Assurance" or "Information Operations", they have maintained the warfare mentality when it comes to overcoming cyber attacks. They frame the problem in terms of an offense and defence. They believe in the value of collecting intelligence data on the potential attackers. The IW warfighter expects that the offense and defence will both engage in surveillance and reconnaissance geared at identifying target systems and threats. They understand the importance of analyzing the impact of a "hit" on mission effectiveness, and conducting immediate damage assessment. This warfare mentality fosters an understanding of the total security picture, assists in identifying where and how to deploy defences, promotes quick damage assessment and early recovery strategies, and leads one to see technology gaps that need to be filled.

Survivability as a Goal

Second, the military's need for survivability is now applied to computing and communications systems, as well as traditional weapon systems and people. Most INFOSEC professionals don't think in terms of "survivability", but that is exactly what information resiliency is about. Defence contractors build systems that respond to adversary IW threats in response to the military's requirement for "survivability". While traditional survivability has focussed on such vulnerabilities as counter-mine for ships, and counter-air for aircraft, more recently we have seen Counter-IW requirements being added to the specifications. In response to this requirement, defence contractors, who are doing their jobs right, should gain a thorough comprehension of the adversary threat (in the context of the platform mission), identify the impact of the adversary threat upon the mission, and identify mitigating technologies, techniques and tactics. The IW warfighter gains a much more complete and realistic impression of the threats to critical information and communication systems and consequently, the steps that he must take not only to protect them, but to ensure the continued operation of his system, and the uninterrupted flow of information to the warfighter, in the face of an IW attack.

Understanding the Cyber Attack Cycle

Third, the concept of information resiliency, and the components that comprise it, were developed in response to the US military's need to understand how a cyber attack unfolds and where to deploy its defensive resources, in a familiar framework. The net effect of gathering that understanding was a significant observation that cyber attacks follow a prescribed series of events *and* that those events are somewhat dependent on the preceding event. For example, before a system can be effectively attacked, the attacker has to accomplish surveillance and then perform an analysis of the target system. The surveillance of the target system provides that knowledge of networks, operating systems, etc. that is then used in the analysis to uncover potential entry points and exploit publicly available vulnerabilities. Clearly, these events show a natural ordering.

Defence-in-Depth

Fourth, the military's use of multiple lines of defence is transferable to cyber war. In any type of conflict, a single mode of defence is unlikely to be effective over time. Attackers get smart to that defensive strategy, and then

alter their means of attack. Thus, the military has moved towards Defence-In-Depth in which it utilizes multiple, overlapping types of defence against cyber attackers. For example, firewalls and proxy servers allow two forms of defence to be placed on a system without undue operational impact. Understanding the cyber attack cycle (previous section) facilitates the selection, placement and timing of defensive deployments.

Know Thy Enemy

Fifth, we borrow the notion of reconnaissance and threat analysis from the military. With sufficient knowledge of the adversary, we can anticipate his actions and take proactive defensive actions. By moving out from behind our entrenched defences, we can perform our own surveillance and reconnaissance, characterize the potential attacker's *modus operandi*, and search for indications of the impending IW attack. By this we mean that we must be looking for trace evidence of the adversary's own surveillance, reconnaissance and intelligence gathering activities (as represented by mapping and access probing). The successful identification of these precursors in advance of the actual attack serves to provide an Early Warning System that effectively eliminates the adversary's element of surprise. Armed with reliable knowledge of our adversary and a reasonable hypothesis about what course of action he is likely to take, the Information Protection Manager can cue intrusion detection and evidence gathering systems, and be poised to implement a response and recovery plan.

Flexible Operational Concepts

Sixth, the US military's plans for twenty-first century warfare incorporates flexible operational concepts that offer the warfighter alternative means of achieving tactical or strategic objectives, if the standard or preferred mode of operations is too risky, or the Rules of Engagement do not permit it. Likewise, in cyber war, one needs flexible and multifaceted concepts of operations that allow critical missions and business processes to continue, perhaps in an alternate or degraded mode, in the face of an IW threat. With advances in forecasting intrusive attacks, we will soon be able to pre-plan our responses to the attack. For example, upon intrusion detection we could do any or all of the following: shunt the attacker to a DMZ area and "fishbowl" him; deflect the attack by shunting him out of our network entirely; transition our own operations to another network, abandoning the threatened network (and leaving the intruder behind) in the process; initiate evidence collection for subsequent post-event analysis, evidence gathering and prosecution; backtrack and attack the intruder at his host site; or identify a target for a retaliatory IW strike. All of the alternative actions are forms of *proactive defence*, which we will discuss later in this paper.

Timing Is Almost Everything

Finally, we borrow from the military the concept of upsetting the adversary's timetable and advancing our own. By throwing the adversary off his plan, we may deter him from future action or postpone his attack. With a reliable knowledge of whom our adversaries are and how they operate, we can deter the attack, change the attack method, and/or minimize damage. We must also use our own defensive tactics at just the right time. It is not sufficient to simply know what to do. We have to know when to deploy specific defensive tactics. Detection of an intrusion after damage has been done is sub-optimal, at best. The right defensive tactics used at the wrong time (either too

early or too late) weakens or totally invalidates the effectiveness of the defensive tactic.

To achieve information resiliency, cyber defence must move from a "reactive" paradigm to a "proactive" one. Proactive defence anticipates an attack, and then uses defensive tactics to respond to an attack as soon as, or before, it penetrates and causes damage. To achieve this goal, one needs to understand where a cyber security event is in an attack cycle, in order to best deploy defences or countermeasures. To this end, we have constructed an IW Timeline to aid in the timely deployment of defences.

The Information Warfare Timeline

If one could take a God's eye view of an information warfare attack, and really examine it in detail, two key points become evident. The first is that there is a precedence and relationship to many of the events that occur in an attack. For example, before a system can be effectively attacked, attack access points need to be established by the attacker. Before the attack entry points can be developed, the attacker must gain access to the system. To successfully gain access, the attacker must avoid being detected by the system firewalls and defences. To gain that access, the attacker must perform surveillance and understand the nuances of the system to be attacked.

Therefore, we can identify and examine generic events and the anticipated sequence of those events of an IW attack. We can then understand what events are attacker-dependent and what reactions are defender-dependent. To begin with, we can construct a visual, or timeline, representation sequencing generic information warfare events. Figure 1 shows a time-sequenced overview of the generic actions that will happen given an intrusive IW attack. (We understand that other categories of attack, such as virus attacks or spamming may not follow this timeline.)

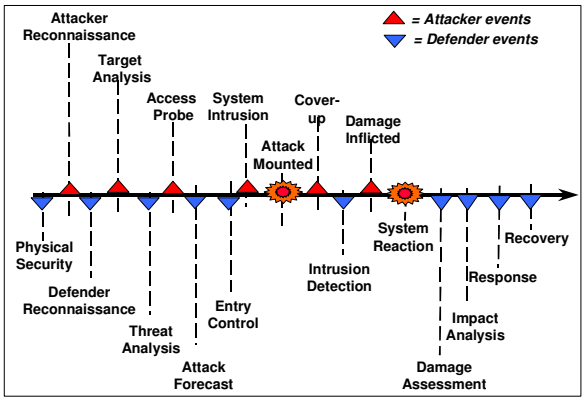

Figure 1: The Information Warfare Timeline facilitates understanding an intrusive attack and defence.

Timeline Events

There are a number of things that need to happen prior to an IW attack and a number of things that must occur after an attack. Each of the events we depict on the IW Timeline in Figure 1 are briefly described below.

- Physical Security refers to the security devices that restrict physical access to computing systems. These include "dumb" devices such as locks, as well as "smart" devices such as facial recognition systems and fingerprint access control (e.g. TrueTouch's Biometric Software Security Suite, IriScan's Iris Recognition Technology).
- Attacker Reconnaissance refers to those actions taken by the attacker to "scout out" a system electronically and through other forms of research. For example, a ping could be done to the system that leads to the mapping of the network and identification of computing and network resources that comprise the system. This

can then be augmented with widely available information on vulnerabilities of the operating system, the router, etc. all of which forms a knowledge base for planning the access portion of the attack.

- Defender Reconnaissance refers to those actions taken by a defender to "scout out" potential attackers and their potential entry points. Some enterprises gather data on potential attackers via "honey pot" web sites designed to lure people interested in attacking entities associated with specific issues. For example, a company or agency engaged in controversial business (e.g. tobacco companies, intelligence agencies, major oil companies during a fuel shortage, political groups) may set up sites with provocative content designed to attract attackers, much like bees to a honey pot. When an unsuspecting person visits a honey pot site, information is taken about the visitor without his knowledge. Such defensive surveillance can be augmented at later points in the timeline as information collected on potential attackers is combined with other anomalous behaviors and observed patterns (i.e., as in a pattern of origin). Information Protection Managers can also use commercial vulnerability scanners (e.g. Internet Security System (ISS)'s Internet Scanner, Netect's HackerShield,) to conduct reconnaissance on their own systems, to determine the weak points that could be exploited by an attacker.
- Target Analysis refers to those actions taken by the attacker that augments the results of the reconnaissance with information available to the attacker on the Internet or through other sources. For example, the reconnaissance could show that the target system is built from Dell boxes, running Windows NT, and using a Novell network. The attacker then can visit several hacker web sites to get known vulnerabilities of these components.
- Threat Analysis refers to those actions taken on the defensive side to further identify potential attackers and their motives. For example, some of the honey pot sites referenced above were found to be military intelligence sites for foreign governments. The "honey" in those cases was supposed to be descriptions of foreign military equipment.
- Access Probe refers to those actions taken that probe the system for access beginning at the fundamental user id/logon all the way to full root access.
- Forecast refers to integration of information by the defenders that allow forecasting of attacks to be made. Typically, this takes the form of information from the reconnaissance and surveillance phases analyzed and synthesized with neural nets and related technologies (e.g. Northrop Grumman's Network Early Warning System).
- Entry Control refers to those measures that are taken to electronically restrict access to a system. These controls may be multiple passwords, card readers integrated into the system control, firewalls, etc. (e.g. Check Point's FireWall-1 4.0, Cisco's PIX Firewall). Entry control devices are also those pieces of software that reside in the router or other portions of the network that automatically allow or deny access.
- System Intrusion refers to the event of actually gaining the access into the system.
- Attack Mounted refers to the actual attack itself, which may occur as a single event or as a wave of attack events.
- Intrusion Detection refers to the synthesis of information that allows the system to realize that an intrusion has

either been completed or is in the process. This is frequently accomplished with commercially available host- or network-based intrusion detection systems (e.g. ISS's RealSecure, Axent's Intruder Alert, or Network Associates, Inc. (NAI)'s CyberCop). The "in-the-process" portion is sometimes referred to as "indications and warnings" where detection is accomplished using reconnaissance and/or access probe information. This is one of those areas on the timeline where the actual knowledge of an intrusion can occur during a period of time that can stretch all the way from the reconnaissance through an attack itself. Obviously, the more sophisticated the detection process, the earlier in the timeline an intrusion can be either detected or predicted

- Cover up refers to those actions taken by the intruder to cover up and eliminate evidence of the intrusion and the attack.
- Damage Inflicted – Refers to the results of the attack in terms of corruption or loss of information and/or system functionality. In any event, this also defines a level of trustworthiness or resiliency that the system still has.
- System Reaction refers to the changes in system operation which result from the attack (e.g. every other blip on the air traffic controller's screen disappears, or the wrong people get billed for overseas phone calls). These reactions could be designed into the attack itself or occur as the result of the system going into a different mode and then failing because of compromised information directly related to that mode.
- Damage Assessment refers to the assessment of the functionality and/or information loss of the system (i.e. electronic battle damage assessment).
- Impact Analysis refers to the defender's projection of the impact of the attack on the business process or mission operations. Ideally, "What if?" impact analyses would be done long before any attack occurs. In that case, procedures and look-up tables would be created to advise the Information Protection Manager about the seriousness of the impact of various attacks on standard operations.
- Response refers to the actions taken by the Information Protection Manager to respond to an attack. Examples include: shutting down parts of the network, eliminating services such as FTP or e-mail from certain network nodes, initiating evidence collection, etc.
- Recovery refers to those actions that are taken to reconstitute the information in a system and return the system to full operational and trusted conditions. The optimum goal of recovery is to non-intrusively restore the system in real-time with no manual interaction.

There are two additional defensive events which should be part of an attack/defence cycle, but which don't lend themselves to depiction at any one point on the IW Timeline: "Security Sensor Integration and "Security Sensor Fusion." Both defensive events occur throughout the timeline, and enhance the Information Protection Manager's situational awareness.

Security Sensor Integration refers to common access to security sensors from a central monitor that combines and presents the output in a unified format. Security Sensor Integration allows the Information Protection Manager to monitor sensors throughout the timeline from one central location. He can see and visualize the output from physical security devices such as badge readers, entry controls such as

firewalls, reconnaissance devices such as vulnerability scanners and intrusion detection systems. The US Air Force's Automated Intrusion Detection Environment (AIDE), and companies such as e-Security, Inc. and Applied Visions, Inc. are pioneering work in this area.

Security Sensor Fusion refers to the analysis and synthesis of the output from a variety of security sensors to correlate security events to each other and formulate a more accurate representation of the state of the enterprise's security. By fusing security information from diverse sources and then presenting the results across the timeline, one gains a view of the entire security state, including what has already happened, what's happening now, and what is likely to happen in the future**.** The area of Security Sensor Fusion is a relatively new technical area, with no widely available tools to support it.

Cyber Defence is Typically Reactive

You will note that in the IW Timeline we just described, there are two distinct chains of events occurring. The first chain is comprised of *proactive* events driven by and controlled by the attacker. The attacker completely controls the time, and can take as much time or as little time as necessary to either cause or effect all the steps in the event chain. The attacker may purposefully space out certain steps over a long period of time so that any chance of detection is dismissed as an anomaly. The other chain of events is *reactive* in nature and describes those actions taken by the system defender.

As can be seen, the attacker has a built-in advantage in the time dimension. Prior to executing an attack, the attacker usually performs extensive surveillance on the system, frequently mapping the network to find the area most conducive to entry. Once the surveillance is completed, the attacker will attempt varying levels of intrusion. Typically, the attacker will start by gaining general system access (by acting as a legitimate user) and then extending this access to the root or system manager level. With this level of access privileges, the attacker can become an authorized user by altering the system access tables. Now the attacker is in a position to enter the system at will and cause an attack when and where he or she chooses. Once the attack is planted, the attacker can remain logged onto the system to watch the attack take shape and garner response information along with real time battle damage effects on the operability of the system. Alternatively or following that, the attacker can remove all traces of entry and withdraw from the system.

Operating in the same space with the attacker's actions, but frequently in a completely different time frame, is the defender's timeline. The defender is almost always in a reactive state, and those actions commence only after an action by the attacker has been detected. And that is the key point, detecting the intrusion. Once an intrusion has been detected, the defender becomes alert to the prospects of an attack and can begin to examine the system for evidence of an impending or on-going attack. More often than not, the defender does not find the attack itself until after the attack is completed. Once the attack has completed, the defender performs damage assessment to ascertain where the damage occurred, what type of damage was caused, and how the attack happened. The defender can then select the optimum strategy to recover compromised information and regain trusted operation of the system.

We believe that the defender should strive to have his reactions occur as close to the

attacker's actions as possible, thereby compressing the timeline and disrupting the timing of the attacker's plan.

The Goal of Resiliency … Compression of the Timeline

To achieve information resiliency, we must drive portions of the defender timeline closer to corresponding events in the attacker timeline, as illustrated in Figure 2. If the defender uses his own reconnaissance and forecasting systems to identify the nature of forthcoming attacks, the defender can then "cue" the intrusion detection system to look for certain behaviors. As a result, the lag time between the actual intrusion and the time it is detected will be compressed. Similarly, the defender can use the results of his own reconnaissance and detection to prepare damage assessment, response and recovery systems for "incoming" attacks. We want to understand the potential scope of the damage before actual damage has occurred. The assessment of damage and the recovery of the system can then commence as soon as the attack is initiated. In fact, certain recovery strategies that require the backup or saving of critical data can be started simply based on the prediction of an attack, before it is actually detected.

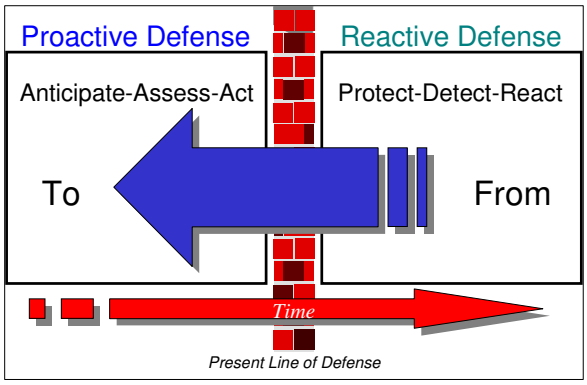

Figure 2: The Defender must move from a reactive to proactive position to achieve information resiliency.

By compressing the timeline, we put cyber defence in a proactive, rather than reactive, role. A proactive defence should be able to:

- Understand the potential threat sources and their *modus operandi*
- Analyze message traffic to predict an attack before it occurs or discover an attack in progress;
- Control the access of authorized users and deny access to unauthorized users
- Integrate the output of various security sensor and fuse the information into a consolidated picture of the security state;
- Assess the "battle damage" in real time;
- Plan for and effect the rapid recovery of any compromised information contained in the system; and
- Maintain event logs that can be used to support subsequent prosecution or selection of counter-measures (offensive IW response).

This may seem to be a daunting set of goals, but they do not have to be implemented all at once. Using the IW Timeline model as a guide, the Information Protection Manager can analyze where he needs to shore up his defences first, and then institute procedures for that area. Later, he can incrementally add features and strength to the system using the timeline as his model.

Summary: "Good Enough" Isn't

Reliance on defensive systems that simply detect intrusions yields few operational advantages to the IW warfighter. Although it may be possible to tighten defences upon the detection of an intrusion, the means through which this is accomplished may well benefit the adversary more than the system to be protected. This is especially true, if "in effect" your defensive response is to self-impose a denial-of-service by cutting

communications connectivity in order to expunge the attacker from the system.

At this point we reach the boundaries of where commercial products offer value to the solution. For many applications, it is simply "good enough" to detect the intrusion and to take whatever means necessary to weather the storm by battening down the security (and access) hatches and shutting down connectivity access (either partially or fully) to the outside world. This "good enough" solution does not support the objectives of information resiliency, as shutting down and restraining access unnecessarily restricts business processes and critical operations.

The Authors

Paul Zavidniak leads Logicon Inc.'s Information Warfare intrusion detection and forecasting activities, focussing on IW survivable communications and the development of advanced warfighting concepts and technologies to meet the military's tactical requirements. The transference of his work in adversary threat of operation development from the military environment to critical infrastructure protection analysis, IW impact assessments, and concepts offers ground-breaking potential. He may be reached at mzavidniak@logicon.com.

Dr. Anita D'Amico leads an Information Security consulting practice at Applied Visions, Inc. (www.avi.com), focussing on IW situational awareness and the transferring of IW technology from military to commercial applications. From 1994-1998 she was the head of Northrop Grumman's Information Warfare business area. Dr. D'Amico's training in experimental psychology provides her with a strong foundation for analyzing IW situational awareness. She can be reached at AnitaD@avi.com.

Dennis H. McCallam leads Logicon Inc.'s Information Assurance and real time information recovery activities, focussing on real time system resiliency along with the development and integration of advanced information assurance technologies to provide layered IW defences. From 1996 through 1998 he was the Chairman, USAF Information Recovery IPT to coordinate and leverage the work being done by several recovery efforts. His unique contributions to information warfare defence, including of the development of minimal data set and half-life concepts, are being applied in a number of US and coalition real time information systems. He may be reached at dennissmcc@home.com.