

**Applying the Concept of Minimal Essential to Maintain
Operational Continuity and Attain Mission Assurance During
Internal and External Attacks on the Information Environment**

Track 7, IS/IO

Dennis H. McCallam
Northrop Grumman Information Technology
1831 Wiehle Avenue, Suite 100
Reston, VA 20190-5241
Telephone 410-925-3223; Facsimile 703-318-9464
dmccallam@msn.com

Perry G. Luzwick
Northrop Grumman Information Technology
1831 Wiehle Avenue, Suite 100
Reston, VA 20190-5241
Telephone 571-203-6131; Facsimile 703-318-9464
pluzwick@northropgrumman.com; pluzwick@comcast.net

Applying the Concept of Minimal Essential to Maintain Operational Continuity and Attain Mission Assurance During Internal and External Attacks on the Information Environment

Dennis H. McCallam

Northrop Grumman Information Technology
1831 Wiehle Avenue, Suite 100
Reston, VA 20190-5241
Telephone 410-925-3223; Facsimile 703-318-9464
dmccallam@msn.com

Perry G. Luzwick

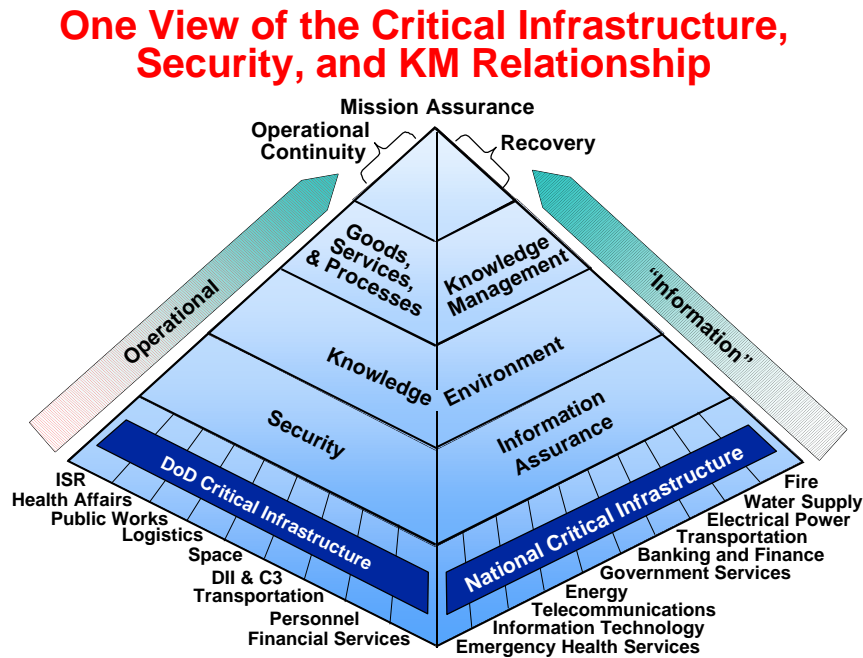
Northrop Grumman Information Technology
1831 Wiehle Avenue, Suite 100
Reston, VA 20190-5241
Telephone 571-203-6131; Facsimile 703-318-9464
pluzwick@northropgrumman.com; pluzwick@comcast.net

Abstract

Operational continuity while under sustained physical and cyber attack is possible using the concept of operational continuity. To achieve operational continuity, in addition to a synchronized and coherent defense of the information environment, data resiliency must also be used. What takes data resiliency from theory to practice is the use of minimum essential data sets on critical information. Through frequent updating of critical information and its dispersion throughout a network, robust sets of essential data can be maintained. Because network centric warfare and information superiority are important in achieving military successes, data resiliency for operational continuity is essential for achieving mission assurance. Data resiliency and operational continuity are important features within Coherent Knowledge-based Operations™. Continued research and operational testing is needed to mature the concept of operational continuity, such as integrated physical and cyber continuity of operations.

During InfoWarCon 2001 in Washington, D.C., a senior Department of Defense (DoD) official said two IT-related DoD organizations were actively managing a virus by judiciously disconnecting portions of the NIPRNET and other .mil systems to prevent spreading the virus, limit damage, and reduce entry points for the virus into the network. After more than eight years of work on protecting the information environment, the defensive measure of choice is still to inflict a self-imposed denial of service. No matter how well intentioned, the result is loss of operational continuity and information superiority, reduced or total inability to conduct network operations (either network centric warfare or network centric business operations), and failure to achieve mission assurance. This is a clear call for employing processes and associated technologies that maintain operational continuity and attain mission assurance during an attack on the information environment.

There are three critical infrastructures that support the information environment: power, telecommunications, and information technology. Unfortunately, many believe the information environment is somehow not associated with critical infrastructures. Degrade, damage, or destroy these infrastructures, and the information environment will become unreliable and then unusable. When the information environment reaches these levels, businesses will have a drop in revenues and profits, the government will not be able to execute functions for its citizens, and the military will be severely constrained in marshalling, deploying, employing, and redeploying its forces. Protecting critical infrastructures is a requirement for achieving mission assurance (figure 1).¹



Northrop Grumman Private / Proprietary Level 1c

Figure 1. A strong foundation is necessary for mission assurance.

Traditionally, attacks have been categorized as external and internal to the system. Examples of externally caused problems are the chemical train fire in Baltimore, the events of 9-11, and the Nimda and Sircam viruses. Examples of internally caused problems are the stock price manipulation in Emulex Corporation and the logic bomb shutting down manufacturing at Omega Corporation.

The new paradigm is that both are forms of insider attacks since an outsider who has logged onto a system (via hacked, stolen, or real credentials) has now become an insider. Once cyber perimeter defenses have been breached, the mission of the assurance portion of the system is to detect the malicious activity and either stop that activity before any damage can occur or prepare for the recovery and reconstitution of the information within the system. The key, and the challenge, is to perform this in real-time, not to discover the attack after it happens when it is far too late to maintain the system's continuity.

Most systems today are composed of open systems, commercial components, remote monitoring, remote access, and standard protocols. These are well known quantities with a host of information available on the Internet. And although you can count on the hardware operating and remaining available, providing any kind of information security to the data in the system is much more challenging. Compounding the problem is the desire to maintain the functioning of the system in real-time; the enterprise, whether business or military, needs to continue the operations with no customer impact. In addition to these shortcomings, much of the previous work on intrusions into all types of systems naively assumes that protection, such as firewalls, is sufficient, thereby making recovery of essential information unnecessary. The hacker intrusion and takeover of the New York Times server on the weekend of 12 September 1998 and the subsequent damage caused, solidly demonstrates the fallacy of too much faith in existing firewalls as a sole source of protection.² As the hacker continues to get smarter, the need to replenish compromised information and to do that in real-time gets more and more challenging. Current data recovery approaches do not take into account the urgency of the data that must be recovered since most all legacy recovery approaches rely on the off-line "cold start" techniques that result in significant system down time. This also often results in the loss of data that existed either at the application level or within the database at the time of the information event. So, this leaves the question, "Is it possible to employ techniques, tactics, and strategies that can help information systems maintain operational continuity and do so in real-time?"

The answer is yes. The Data Resiliency in Information Warfare (DRIW) Program, an effort sponsored by the US Air Force Research Laboratory (AFRL) and developed by a Northrop Grumman Information Technology, formerly Logicon, team³, demonstrated this form of recovery and integrated this over an existing system.

This paper will discuss several issues and the process and thinking that would support implementation. In addition, we will present the methodology for how the recovery can operate and then we will demonstrate that recovery on a real system. The presentation will specifically address the following:

- What is meant by real-time enterprise continuity
- The steps attackers take to hack into your system and where the defensive measures currently succeed and fail
- Key technologies for real-time insider detection and real-time system recovery
- Process considerations for implementing system recovery along with an understanding of return on investment for implementing
- Some streamlining ideas and results from recovery implementation


What is Continuity for the Enterprise?

Enterprise continuity can be summed up as the ability for a given enterprise, or information system to continue operation in the face of potential disruption. In the context of this analysis, we consider that the goal is to maintain the continuity of the information system. The ability of the enterprise to respond as a whole is factored in where appropriate, but the crucial element remains the information system. A wide range of events and incidents, both physical and cyber, can adversely affect the operation of the information system. And it is our handling of those

incidents and events that affects perception of the institutions business operations. For example, snowstorms, electrical outages, and hacking attacks must be equally addressed. Our goal is simple, maintain continuity of operations without sacrificing performance and ensure there are no adverse interruptions in the service. In short, recover and continue seamlessly in real-time.

Continuity implies that application of defensive measures without adversely affecting the performance of the system. It does little good to employ defensive measures that eat into the system overhead budget to the point that the applications programs are impacted. The point of the protection measures is to ensure that the system and its information survive any attack. Furthermore, the survival of the system must be such that its integrity and trustworthiness is preserved. It is well known in the security domain that no one tool will provide complete protection against attacks. The concept of layering of defenses emerged where several solutions or tools are used and these cooperate with each other in terms of sharing information so that weaknesses are compensated for. Figure 2 shows span of integrated defense; it involves a number of different areas all cooperating. There is a huge benefit to this in terms of avoiding single points of defensive failure and allows for some creative approaches to be employed. One of the challenges that computer security people face is the collection of cyber evidence resulting from either computer misuse or some direct computer attack. Evidence is now reconstructed using the state of the system in the aftermath of the attack. Cooperative defensive systems operate in real-time and offer the potential to gather evidence in real-time to be able to completely reconstruct the tampering and not omit any key evidence. This is made possible by employing indications and warnings that look for specific forms of pre-attack, shares those warnings by triggering portions of the defensive suite, and helps the system adapt to the environment rather than react to the environment. This adaptability is crucial because the techniques and tactics of the attackers move at a rapid pace. And this is best exemplified by looking at reported and exploited vulnerabilities that accompany the release of a new operating system. Within hours of release, the holes are found and used while it takes much longer to affect repairs.

Some Areas of a Complex Knowledge Environment Defense



Northrop Grumman Private / Proprietary Level 1c


Figure 2. Information assurance encompasses a wide range of technologies and areas.

There is an interesting relationship between the type of event that causes interruption and the reactions that we as either users or just observers have. These reactions are simply the patience shown in dealing with a system or process that has been interrupted. The most patience is usually shown if the interruption is caused by an act of terror or as the result of an act of war. Acceptance of the loss and inconvenience seems to be proportional to proximity of the act. This is supported by the patience and understanding shown to businesses that were directly affected in the recent attack on the World Trade Center. This patience can last for quite a long time. The next event is the natural disaster kind, such as tornadoes, blizzards, and hurricanes. Understanding that an order cannot be fulfilled due to snow blocking the road is accepted, at least for a period of time in the 24 - 72 hour range. By that time, the feeling (right or wrong) is that the snow should have been cleared enough to resume normal operations. Now suppose you are sitting in your family room, watching a favorite program when the power goes out. One of the first things typically done is to look outside and check the condition of the remainder of the neighborhood. If there are other houses experiencing the same outage, then we feel a little (not too much) better about the situation. This seems to follow the old proverb "Misery loves company." When dealing in a more one on one situation involving human error, the patience level drops. Consider the example of being in a checkout line and having the clerk make a mistake that causes the entire order to be re-processed. The reactions to this are varied and are usually the result of an instant judgment. Our patience will increase if the clerk has an efficient recovery and will decay if the clerk fumbles around. Clearly the least amount of patience we have is with a cyber outage, such as the computer 'freezing' or having the blue screen appear.

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 defenses. 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 3 shows a time-sequenced overview of the generic actions that will happen given an intrusive IW attack. There are categories of attack, such as virus attacks or spamming that may not follow this timeline, but in general this is representative of the steps an attacker will go through in order to gain entry to a system. The attacker actions are on the top of the timeline and the defender actions are on the bottom. Understanding the timeline is essential to developing recovery and continuity strategies.



Northrop Grumman Private / Proprietary Level 1c

Figure 3. The information warfare timeline showing both attacker and defender operations.

The Key Technologies

The timeline really presents a good foundation for developing defensive strategies. Consider the example of a perimeter control system. Authorized personnel have either password or access cards that used with the control device, allows entry into a guarded area. These systems are quite

commonplace, and usually records log of entry and associate those entries to authorized users along with maintaining a time record of use. This device is not sufficient to guarantee that the user is who they purport to be, but it could provide some evidence or even some warning particularly if we consider the fusion of information.

Fusion offers significant promise for information assurance. One of fusion's strong points is the ability to correlate associated events and then discriminate between like events. This then allows time synchronized and normalized event data to be compared against the logical and progressive stages of an attack. This then allows the formulation of a set of beliefs based on the observed and collected evidence. If possible, in the preferred integrated defensive suite, other sensors could be tasked to provide corroborating information that could provide verification, validation, and refinement of belief. The bottom line is the expansion of the available knowledge domain to look at and evaluate suspicions.

Another key foundational technology is embodied in the concepts of minimal essential and data half-life. The central issue in information system recovery is to identify what data to restore. Minimal essential data sets (MEDS) are defined to be the smallest collection of data and information that describes a given system and can be used to replenish compromised data in a system. The MEDS approach uses the components of the system to select the ranking of processes and then uses the data descriptions of those processes to develop the minimal sets. MEDS doesn't care as much about the value of a piece of data, rather it cares about the relationships that a piece of data has. Those relationships if taken in the aggregate then define the transformation computations for all data in the system. MEDS, as a concept, has great similarities to the mathematical concept of basis vectors. Basis vectors are by definition the smallest set of linearly independent vectors that, taken in some combination, completely span the vector space. It is a similar relationship that MEDS exploits to span the data space of a given system.

The half-life concept resulted because data elements are computed or refreshed at various timing intervals. In examining the range of intervals, it was found that each data element in the system had a "time of validity" associated with it that would influence its prioritization in a recovery plan. Half-life value of data elements could be used to determine recovery priorities which in turn specifies the elements identified as MEDS, and how often they need to be saved or backed-up within the system.

One of the first application domains for the fusion was the Network Early Warning System (NEWS), a program performed by Northrop Grumman Information Technology for the USAF. NEWS had the goal of identifying potential large scale attacks by correlating anomalies. This was a difficult problem to solve, namely the identification of impending large scale attacks by looking and evaluating potential precursors. This technology provides an alternate approach to just intrusion detection systems, which only work after the intrusion has happened, by correlating events to the timeline and looking for activity which occurs before an attack. The major elements of the NEWS solution are:

- Profiling/Correlation/Clustering - which reduces intrusion detection system clutter using features from incident profiling and associates related attack data together

- Fusion - works to resolve the gaps in heterogeneous attack reporting data
- Forecasting - provides cues for feedback and courses of action based on timeline associations and beliefs
- Visualization - provides analysts with the capability to comprehend critical resource impact by looking at a 'picture' of what is going on⁴.
- By identifying precursors of attacks early, the process can separate large-scale attacks from innocuous or random activity
- The techniques employed enable multi-site correlation and fusion of attacks to provide protection not only for a given single network, but also to networks of networks.

Fusion was employed for a different purpose in the detection of an insider attack as the attack is progressing. IAMPS, the Insider Anomaly Measurement Processing System, looks at the authentication consistencies and the authentication inconsistencies to ascertain that a given person is in fact, who they claim to be. Consider the case of a valid user ID/password combination, a token for access into other parts of the system, devices (such as card readers) that allow physical access to computer equipment and information from the timekeeping system. Suppose the following case is considered: a Joe ID - Joe Password are valid users of the system, Joe Access Card, and Joe's record of his time for this week. One form of consistency is that Joe used his access card to gain entry to the area, then he used his Joe ID and Password to log on to the system the system, and Joe has time entered on his timecard for work on Project X. Taken one at a time, each of the data items are valid and would indicate that the person logged onto the system is in fact, Joe. This is an authentication consistency. But suppose the time card indicates Joe is on vacation. Again, each of the individual data items are valid and if taken one at a time, would indicate Joe could be on the system. But taken as a unit, they show an authentication inconsistency in Joe actually being on-site. At this point, there are courses of action that can be taken, the least of which would be to terminate the access from 'Joe' until a follow-up can be completed. In essence, we are developing a process where we are trying to prove that someone logged in to a computer is either who they claim to be or is impersonating a legitimate user. In fact what we are trying to accomplish is to view the data, correlate where appropriate and present evidence.

Original Recovery Application Domain

Recovery technology built on the MEDS and half-life concepts and added some granularity on types of system recovery. Essentially, there are three defined types of recovery:

- Hot - Where the system is fully recovered and reconstituted without any operator guidance or assistance. In some cases, the event is missed by the operator and thus preserves real-time mission continuity.
- Warm - System is partially reliant on 'human in the loop' to complete recovery and reconstitution. ("system has an exception at.....it may be possible to continue normally")
- Cold - System requires a complete re-boot (the blue screen)

In the development of the recovery technologies, a target system was selected to provide a real world foundation and a realistic architecture that would channel and guide the solution. The system set that was selected came from the airspace management family of systems. Airspace

management systems provide command and control for aircraft operating in a given airspace. As such, there are two types of systems, air traffic control or those systems that provide the command and control for civilian or commercial air traffic, and air defense systems that provide the same control for military aircraft. The processing, interfaces, and applications software between air traffic control and air defense are almost completely similar except for one distinction. Air traffic control seeks to provide separation between aircraft for safety reasons while air defense systems seek to bring aircraft closer together for military reasons.

But the airspace management systems have a wide range of components that span potential recovery issues. These components are: communications with outside systems, a range of inputs from other systems that impacts internal computations, distributed network, several databases, operator interfaces, sensor inputs, etc. all of which represent portions of most every other real world system. Selecting a system with these attributes would help provide a solid foundation for technology transfer to other systems.

Architecture of the Solution

Implementing a recovery solution requires some modifications to the traditional architecture of an information system. In essence, the recovery applications software is an adjunct to the system and can either be embedded in the system or be a separate component. The key issue is to be able to invoke the recovery execution from the system. By integrating with other information assurance components, the recovery system can maintain a proactive stance in regards to being ready to recover. Recovery is a series of processes that occurs in three distinct phases with tasks that are performed: in expectation of recovery, immediately after the attack is detected, and during data and information reconstitution.

During pre-IW attack conditions the recovery system will:

- * Execute at pre-determined interval (inside of base timing interval) to catch all computations occurring between major data and information updating
- * Perform the MEDS selection and parameter updating
- * Perform the MEDS compression and encryption
- * Perform the MEDS storage
- * Be cognizant of maintaining three times - time of last update of data and information data, time of last MEDS selection, time of last MEDS storage (might be slightly different, but important for any information extrapolation purposes)
- * Continue to clear recovery flags to prevent system from "lapsing" into recovery when no recovery is warranted

Immediately after the attack is detected, control initiates the recovery activities that will:

- * Set reconstitution hot recovery flags
- * Ascertain damage and focus of attack
- * Select recovery routines to execute based on location of attack; for example, is this an instruction memory re-load or data memory reconstitution?
- * Invoke special re-start up routines to restore constant constants, if necessary

- * Access and de-encrypt, decompress MEDS
- * Restore MEDS

The intent of this portion of the recovery process is to ensure that the system continues operation. The final phase of recovery occurs during the reconstitution of the compromised information where:

- * Interconnections to portions of the system or other systems are re-established by informing connected systems of recovery in progress
- * Secondary reconstitution, if necessary, is initiated, and finally
- * Recovery flags are cleared and the system is returned to normal system operation

Interdependencies and the Domino Effect

Connections to other systems impact recovery in several ways. Recovery uses a variety of means to replenish the information that may have been compromised and some of those methods depend on retrieving minimal information sets from other portions of the system/network. When the MEDS were selected during the pre-recovery phases, certain assumptions were made. Some of the information to be replenished has some computational dependencies on MEDS and relies on the ability of the system to access the MEDS and use those computational dependencies to complete the data reconstitution process. There are many other interdependencies; power and communications are good examples. Here, the situation could be more acute. If, for example, there is reliance on a third party provider for ISP communications, then basing a full recovery without considering any limitations they might have could have an adverse affect. With the goal being a real-time recovery that maintains the operational continuity of the system, understanding the interconnections and the impact of their recovery issues is important.

In short, there are several interdependencies that should be reviewed and analyzed. First is the functional connectivity looking at all the subsystems directly connected to the system. This involves not only the data flow between these systems, but also the availability of those systems to provide any MEDS or corroborating data that would be used in the reconstitution process. Second is to have knowledge of the network configuration. Knowing what is connected and what should be connected allows a building block approach to take place in restoring full operational connectivity. Third is to look at the data computational dependencies between sub components of the system. In the restoration of data process that occurs during recovery, knowing which piece of information may reside where is crucial. The fourth area is to look at the database structure and how information is stored in both the logical and physical sense. Restoration of the information in a database can be a lengthy process and impact the real-time goal of the recovery process.

Recovery ROI

Quantifying the true return on investment is a difficult task. In some respects, recovery has been thought of as an investment in avoidance. That comment is not accurate since recovery is the primary means to effect operational continuity on a given system or across an entire enterprise. The best way to measure ROI is to compute and consider the cost of down time to a given

system. Most operations can be measured in terms of money per time unit for sales. And this applies to enterprises that rely heavily on e-business for sales and those that use information systems in a more traditional manner. In either event, there can be boundaries on the potential lost revenue computed. Presenting accurate ROI has been a subject of great debate in the technical community for years. Often overlooked is the fact that information assurance is very much an issue of risk avoidance, making agreement on calculating ROI even more difficult. One way of looking at it is there is time required to get a system back on line and that time is all related to cost of operating the business. Part of the impact is in sales, because without the system, orders cannot be processed. Given that, the following is a way of calculating the return on investing in defensive and continuity technologies.

The amount of investment at risk can be estimated as:

$$(Cob + Sd * Rf) * Td$$

Where:

Cob = Cost of Operating Business 1 day

Td = Elapsed time to regenerate information (in days)

Sd = Daily sales

Rf = Recovery factor (nominally 2, since it takes twice as long as you think to get back on line)

Consider the following example of a company doing \$10M/yr in sales operating at 10% profit.

Cob = \$24,660

Td = 2 weeks

Sd = \$27,400

Rf = 2

So the at risk investment of not doing any information assurance can be estimated to be \$1,112,440.

Streamlining the Continuity

Once the plans and processes for the recovery and continuity have been implemented and integrated, there begins a constant re-evaluation of the plans and processes to factor in changes in the methods of cyber attacks. Initially, the continuity planning needs to consider a wide range of faults. This range includes the more traditional natural disasters, acts of terrorism or acts of war. These events are far more understood than the cyber events that may not show up until long after the attack was actually completed. Integrating various forms of cyber protection and linking that as 'triggers' to recovery provides a seamless approach for affecting the real-time continuity. Constant re-evaluation of the entire defensive posture ensures that the system will remain protection current, which ensures that the latest cyber countermeasures are put in place. Conversely, the ability to maintain awareness of the latest vulnerabilities factors for a robust defense is needed.

The Demonstration System

The Data Resiliency in Information Warfare (DRIW) Capability Demonstration effort successfully combined and transferred technology from three USAF Study programs through a demonstration of hot information recovery using intelligent agents on a real battle management command and control system. The program, sponsored by the AFRL in Rome, NY and performed by Northrop Grumman Information Technology, incorporated and extended the work conducted on all three programs.

The DRIW program focused on the recovery and response phase of an information warfare (IW) attack and was a natural progression from the previous research efforts. DRIW provided for the first time, an implementation of recovery concepts on a real deployed system. From the outset the effort sought to demonstrate two concepts that proved information resiliency is a valid concept and that intrusion tolerance is achievable. The first concept is that a hot restart of a real-time Battle Management System (BMS) can be conducted in real-time to provide full information recovery and reconstitution. The second and equally important concept is a marked reduction of system recovery time for an attacked system as depicted in the IW timeline. The demonstration system, as mentioned previously, was the Air Tasking Order portion of a command and control battle management system. A hypothetical scenario based on a hostile invasion, was developed and is illustrated in figure 4.

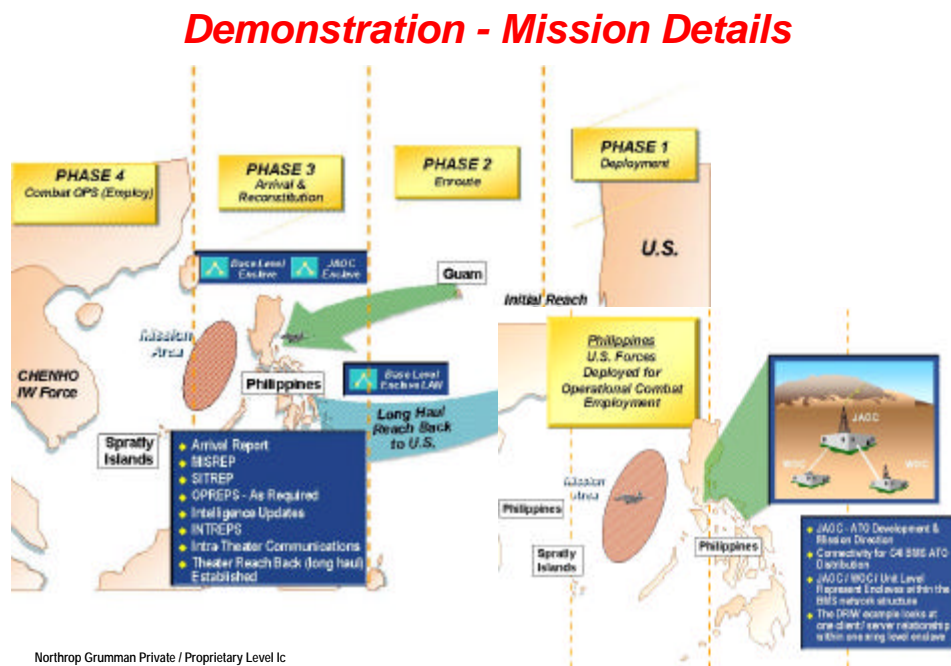


Figure 4. Hypothetical scenario showing all phases of the operation.

In the scenario, the blue forces were moving aircraft to support the Philippines. During the course of moving the aircraft, a refueling mission was scheduled. During the refueling mission the red, or Chenho, forces hacked into the air tasking order system and modified the intercept points. In the demonstration, we performed three hacks into the system and then measured the


system in terms of operational continuity and response to the hacks. In all three cases, system functionality was restored along with the data without any loss of continuity. From the research, implementation and resultant demonstration in this program, DRIW proved that the application of advanced real-time techniques in the recovery of information enhances the performance and operational integrity of (specifically) C2 battle management systems. The two key tangible results of the DRIW are:

1. Insertion of protective and recoverable technologies into current critical infrastructure systems and legacy systems is possible. This becomes an extremely important consideration to improving the information assurance features of legacy systems. Specifically, DRIW shows that real-time recovery can be implemented on legacy systems and real-time command and control systems.
2. The overhead required to maintain the level of protection required for information and system reconstitution was shown to be for the most part, below 12 percent and always within acceptable timing windows. In short, the recovery process does not interfere with the real-time nature of the system. One of the original concerns in the implementation of any recovery system was the potential for additional processing overheads. Without any attempt at streamlining or optimizing the recovery routines, reconstitution times were well within system processing tolerances.

Full Protection

Clearly there are technologies and implementations that will help the defensive posture of an information system. Figure 5 relates the value of full dimensional protection and the graphic makes sense if you think about the ramifications.

The Value of Full Dimensional Protection




- No protection provides little, if any, resiliency
- Static approaches utilize single point or federated solutions
- Integrated approach shares information across sensors
- Increase in operational availability and continuity while providing adaptive approach to IA problem

Figure 5. Relating the value of full dimensional protection.

No protection provides little, if any, resiliency and in fact does not take much of an attack to begin to degrade the system in terms of fidelity and performance. The more traditional approaches utilize single point solutions and are static in their implementations. These solutions if applied in depth are federated at best. They do provide some protection, but in the long run will degrade more linearly. The integrated approach uses the same layers but shares information across and between layers. This defensive suite is now adaptive as opposed to reactive. The impact is an increase in operational availability and continuity while providing adaptive approach to the IA problem. Coupled with real-time recovery, the system becomes even more resilient because it adapts defensively and reconstitutes when necessary. Clearly, even this level of protection will degrade at some point set of critical data, but the level of attack would have to be fairly severe.

Coherent Knowledge-based Operations™ (CKO™)

Operational continuity enables organizations to continue essential operations in the face of sustained physical and cyber attacks, either on the battlefield or in the marketplace (there are many similarities, such as leadership, intelligence, maneuvering forces, and surprise), to attain and maintain competitive advantage and to be successful (either victories or profits).⁵ Operational continuity cuts across all organizational functions. A process to enhance operational continuity is Coherent Knowledge-based Operations™ (Figure 6). CKO™ is the combination of information operations (IO), knowledge management (KM), and Network Centric Business™ (NCB™). Combining operational continuity with CKO™ has powerful advantages.



Northrop Grumman Private / Proprietary Level 1c

Figure 6. Coherent Knowledge-based Operations™ coherently combines capabilities for synergistic enterprise-wide advantages.

At the intersection of NCB™ and IO are tactics, techniques, and procedures to seize market share. One of these techniques is operational continuity. To implement optimum operational continuity requires situation awareness and information superiority. Situation awareness is the intersection of NCB™ and KM. KM enables enterprise-wide expertise to blend with the business strategies and methods of NCB™. Information superiority is the result of combining KM with IO. Using capabilities to attack and defend the knowledge environment must be carefully coordinated using the best knowledge in the enterprise. The powerful advantages of Coherent Knowledge-based Operation™ make operational continuity a proactive resource.

Conclusion

Operational continuity while under sustained physical and cyber attack is possible. To achieve operational continuity, in addition to a synchronized and coherent defense of the information environment, data resiliency must also be used. What takes data resiliency from theory to practice is the use of minimum essential data sets on critical information. Through frequent updating of critical information and its dispersion throughout a network, robust sets of essential data can be maintained. Because network centric warfare and information superiority are important in achieving military successes, data resiliency for operational continuity is essential for achieving mission assurance. Data resiliency and operational continuity are important features within Coherent Knowledge-based Operations™. Continued research and operational testing, such as integrated physical and cyber continuity of operations, is needed to mature this concept.

References

¹ “Mission Assurance. Is Your Foundation Strong Enough to Assure It?,” Computer Fraud and Security, Elsevier Science, March 2002.

² Security expert explains New York Times site break in. CNN, Sci-Tech, By Ellen Messmer, September 18, 1998.

³ The team consisted of Logicon, Inc. (now Northrop Grumman Information Technology), Modus Operandi, Inc., and George Mason University

⁴ Perhaps the best example of visualization occurred in the movie Jurassic Park when the systems were brought back on line. The visual for that sequence in the movie left no doubt as to what was occurring and in what order systems were restored.

⁵ “Gaining Competitive Advantage and Protecting the Information Environment by Using the Principles of War and Knowledge Management,” Lt Col Perry Luzwick, US Air Force, Information Security Bulletin, July 1999.

Rapid Recovery of Information for Real-time Intruded Systems, United States Air Force Technical Report #AFRL-IF-RS-TR-2001-55, Dennis H. McCallam, Cathy K. Piggot, Ronald E. Newland, April 2001.

“Achieving Information Resiliency,” Elsevier Science, Information Security Technical Report on Information Warfare, Vol. 4, No 3; London, United Kingdom, August 1999.

About the Authors

Dennis H. McCallam is on the Senior Technical Staff for Northrop Grumman Information Technology. He leads Northrop Grumman'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 defenses. He is currently the Technical Director for two Information Assurance programs and consults across a number of other information assurance activities and programs. Prior to that he was the Technical Director for the USAF Rapid Recovery Program and the Defense Engineering and Research Agency Recovery Program, Phase 1. His unique contributions to information warfare defense, 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 has three patent disclosures in information assurance and has delivered numerous papers at Information Warfare symposia both in the US and overseas. Mr. McCallam has 27 years experience in the development and applications of real-time software.

Perry Luzwick is Director, Information Assurance Architectures at Northrop Grumman Information Technology. He is a senior technical consultant throughout the corporation for Information Warfare (IW), Information Assurance (IA), Information Superiority, Critical Infrastructure Protection (CIP), and Knowledge Management (KM) projects from conceptualization through design and implementation. He retired from the United States Air Force as a lieutenant colonel. While on active duty, he served as Military Assistant to the Principal Deputy Assistant Secretary of Defense for Command, Control, Communications, and Intelligence. Other assignments included serving at Defense Information Systems Agency, the Joint Staff, National Security Agency, and Supreme Headquarters Allied Powers Europe. He earned an MA, and was a Distinguished Graduate, in Computer Resources Management from Webster University; an MBA from the University of North Dakota; and a BS, Psychology from Loyola University of Chicago. In January 2001 he began his doctoral studies in Knowledge Management at George Washington University. He has taught as an Adjunct Faculty for the University of Maryland, NSA's National Cryptologic School, and the City Colleges of Chicago. Perry has published 20 articles and lectures on IW, IA, KM, and CIP, writes the "Surviving Information Warfare" column for *Computer Fraud and Security*, and is a 1998 member of the International Who's Who of Information Technology.