

“Help! I’ve been attacked!”
**Researching Ways to Recover a Command and Control System Following an
Information Warfare Attack**

Mr. Colin L. Nash*

Strategic and Tactical Systems
Infosec Business Group
Defence Evaluation and Research Agency
St Andrew's Road
Malvern
Worcestershire
WR14 3PS UK
Phone: +44(0)1684 895563
c.nash@eris.dera.gov.uk

Mrs. Cathy K. Piggott

Northrop Grumman Corporation
Electronic Sensors and Systems Sector (ES³)
PO Box 1897 – MS 983
Baltimore, MD 21203
USA
Phone: 410-765-9305
cathy_k_piggott@md.northgrum.com

Abstract

In 1997, the United States Air Force (USAF) Rome Laboratory (RL) initiated research into techniques and methods for recovering from an Information Warfare (IW) attack. Prior to this time, work in the area of recovery was rare if not non-existent. Simply put, the way to recover from an IW attack was to treat the system as if it were coming on line for the first time. This is referred to as a “cold start.” In a real-time Command and Control (C2) system, a cold start recovery process poses a serious impact on the total defense picture. The commander, operators, and users must be able to count on this system at all times to provide information and real-time situational awareness. To avoid relying on cold start recovery processes, Northrop Grumman researched ways to rapidly recover information within an intruded system. Countermeasures were developed for the weaknesses portrayed in the current cold start methods of recovery. During 1998, the UK Defence Evaluation Research Agency (DERA) extended the work done under the USAF activity. DERA initiated an extramural contract with Northrop Grumman to examine the emerging CIS BMS architecture and postulate concepts for recovery from an IW attack. This DERA-sponsored activity transferred and broke new ground on recoverability for distributed communication systems.

1.0 Results

During this eight-month programme, DERA and Northrop Grumman made significant progress in defining a leading-edge recovery scheme for the CIS BMS. The key results can be summarized as follows:

- Established a *multiple timeline taxonomy* that examined the effects of multiple attacks against multiple areas in the system. This concept provides for a visualisation of the attack sequencing.
- Expanded the concept of *Minimal Essential Data Sets (MEDS)* to distributed communications systems, infrastructure, and security. This extended analysis began with a look at how business processes spawned basis sets of data, communications, network requirements, security, messages, and control. These sets were then evaluated and an overall process for establishing the MEDS was defined.
- Extended the concept of the *Half-Life of Information* as it relates to system recoverability. The research looked at the process for and some mechanisms into preserving information within a MEDS based on half-life.
- Developed a process that defines how to establish *the relative recovery priorities of the individual Key Business Functions* based on conflict type and scenario phases.

Each of the above areas is discussed in greater detail in the following sections.

2.0 Multiple Timeline Taxonomy

During the conduct of this research, it became apparent that the methods of attacking a multiple user network were different from the approaches that are normally used against a single system. Simply put, multiple user systems invite more forms of attack. Through our research, we identified and examined the generic events in the anticipated sequence of events or stages of an IW attack. This was extended to look at what events are attacker-dependent and what reactions are defender-dependent. Once the analysis of the attacker and defender timelines was understood, we applied the timeline combination to a multi-user network to see the effects of such an attack. Figure 1 shows a combined timeline considering the actions of both attacker and defender. The following paragraphs explain the actions of each actor.

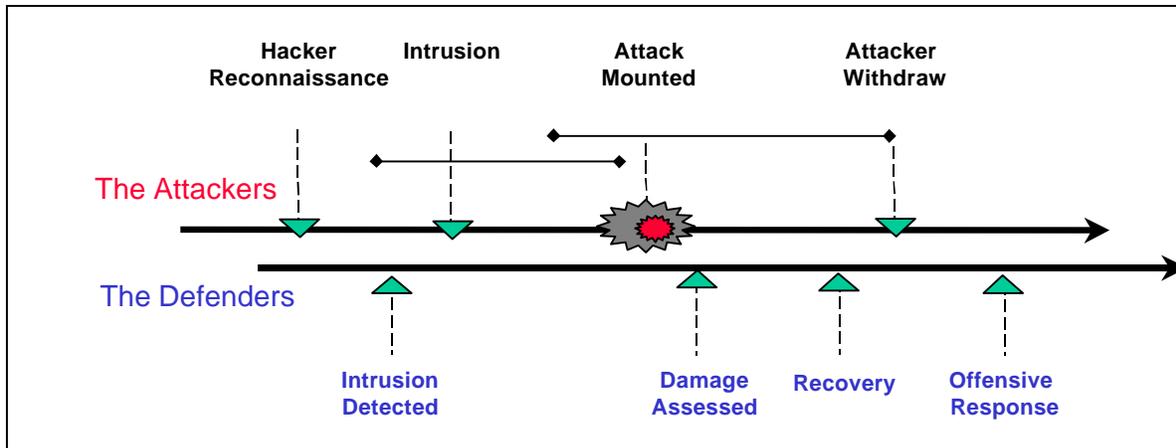


Figure 1. The Attacker and Defender Operate on Different Timelines

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 authorised 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.

Coincident with the attacker's actions is the defender's timeline. The defender is usually in a reactive state, with any positive actions commencing only after an action by the attacker has been detected. 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. There is, however, the possibility that an intrusion can be detected and stopped before the attacker can cause any damage. 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 operation of the system. Another output of the assessment is to select and initiate any offensive responses to the attacker.

The key objective in analysing and understanding the dual nature of the timelines is to drive portions of the defender timeline closer to corresponding events in the attacker timeline. For example, the optimum result from intrusion detection is to acknowledge an intended attack by analysis during the attackers surveillance phase. Similarly, the assessment of the damage and the recovery should commence as soon as the attack is initiated. Given fully interoperable defence-in-depth components, the intrusion detection tools will identify a pending attack and remove any

threat from effecting such an attack. In addition, the attacker identification and location evidence gathered by the defence-in-depth components can be used to initiate an offensive response in some form. This is shown in Figure 2, which highlights the overall goal of shrinking the time lags between attacker and defender. This approach changes the modus operandi of the defender from solely reactive to a more proactive stance.

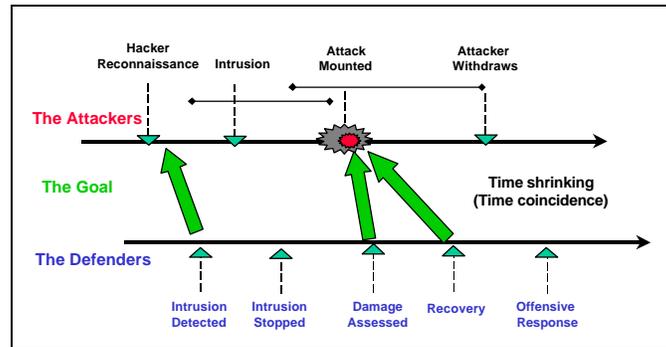


Figure 2. Moving the Defender from Reactive to Proactive

3.0 Minimal Essential Information Components

In any system, particularly a mathematically-based system, there exists a set of basis vectors that defines and spans the system space. Appropriate combinations of the basis vectors then define all the members of that space. Applying that concept to electronic information systems, there would exist a set of data/information that acts as the “basis vectors” for a given system. Given this set of data/information, any other information component in the system can be computed by utilising the basis information. We call this set *the minimal essential information set* and it defines the smallest set of data/information required to replenish/reconstitute the system.

The methodology for selecting the elements of the minimal essential information set is based more on characteristics of the information rather than on the value of the information itself. In essence, there is a hierarchy of computations that describes any piece of data/information in the system. Viewing this, a person can establish a precedence in terms of computations between data elements of the system, describing the “relative importance” of a data element. This relative importance will then lead to the prioritisation of the recovery of that information.

There are several approaches that can be taken to evaluate the precedence:

- Examine the data for time of creation. It is important to understand when the data is initially created -- in terms of compile time, start-up/initialisation time, etc.
- Data in the system is also examined for specifics on the refresh rate of that information. Some information has longer periods of time between updates (such as date), shorter intervals (such as target position), etc. Knowledge of the refresh intervals becomes important when the period of validity of that information is considered.
- Data can be influenced at random intervals (such as weather updates) and can also describe the parameters for displaying that information.

3.1 *Analysis of Minimal Information*

We then applied the minimal essential information concept to the CIS. We began by analysing the overall information flow of the CIS. The initial thing noted was that each Business Process spawns sets of data, communications, network requirements, security, messages, and control. Each piece of data in the database is related to a process through creation of that data, the modification of that data, and the use of that data. Each process uses or exploits the communications and control to execute its functions. Each process uses messages to communicate between users. In essence, each Business Process has a basis set of information and service requests that fully describe that process. This basis set describes the specific relationships that exist internal to a process for use of services and data. In addition to the basis set of information for each business process, there will be some portions common to two or more processes. This implies that there is some common data that exists across multiple processes that, at least partially, is part of the minimal essential information set of several processes.

By identifying this common set of information, we can begin to describe the shortest path to the recovery of full system operation. The minimal essential information set for each process and the minimal essential information set that is common to more than one process will fully describe the system in terms of data hierarchy. This understanding of the data hierarchy is critical. The next step is to apply the data hierarchy concept to messages and services to describe the “basis set” of services required to describe a process in terms of the communications, network, and messaging use. That in turn defines the priority for the recovery of services and other system features.

3.2 *Minimal Essential Data Sets*

Minimal essential data sets (MEDS) are the minimal collection of information that describes a given system and can be used to replenish the compromised data in a system. Selecting the MEDS requires that the system be fully defined and the data dictionary for that system be available for MEDS selections. 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. The process of decomposition uses both the functional breakdown and the data breakdown. With the evolving state of the requirements for CIS, the complete data dictionary is not available, so we applied the concept at a higher level. In this study, we examined the communications services, messaging, business processes, and system level information. Using this information, we then developed some observations on how to effect improved recovery.

4.0 Half-Life of Information

All data within a given system has a time window for both validity and value. For example, the constants “one” (1) and “pi” (3.1415...) are valid at all times. Given this concept, the following can be shown:

- Individual pieces of data can be tagged to indicate the real time of last computation.
- Data can be changed at a periodic time or periodic logical rate based on several conditions being met.
- Data can also be integrated over a period of time. This usually involves filtering (e.g., Kalman Filtering) of the data.

Therefore, in a real time system, each piece of data has the equivalent of an atomic half-life. Half-life is defined as the time required for half of the radioactive atoms (or in this case, pieces of data) in a sample of radioactive material (or in this case, a system) to decay. In nature, the half-lives of radioactive materials range from seconds to billions of years. Continuing the half-life concept to a networked computer environment, this means different pieces of information decay or become invalid at different times. If a piece of data is used in a computation beyond the “time of validity” window, the resultant computation will have induced errors. Additionally, since the CIS BMS is a messaging system, there will be certain messages that become obsolete over a period of time. For example, a message indicating an attack has begun is obsolete at almost any time thereafter. This has security implications, as exemplified by the same message. The command to attack might have the highest security classification prior to the attack; but after the attack has commenced, the classification is overcome by events.

We extended the use of the half-life concept to the development of recovery requirements. For the CIS BMS, this consisted of viewing not only the data resident at the user levels, but also the data resident at the system level, messaging data, information across the CIBIS definitions, and eventually (after full BISA definitions are made) pertinent information across the BISAs. By examining the half-life value for each piece of information in a system, one can start determining the potential recoverability of that piece of information.

In conjunction with the half-life concept is the concept that different data types have different half-life values. As a result, we need to examine each piece of information for its type and its half-life. Identification as to the types of data in a system is necessary to fully develop this aspect of information system recovery.

5.0 Key Priorities

To support the identification and prioritisation process of the suite of business processes (also referred to as Key Business Functions (KBFs) and Battlefield Information System Applications (BISAs)) of the Land Tactical CIS BMS and move towards the goal of information recovery for the business processes, we investigated two areas. First, we analysed the ways in which the stages of conflict could affect the execution of the recovery process. Second, we used the current definitions of the key business processes to begin building a relative prioritisation for application to the recovery process.

5.1 *Stages of Conflict Analysis*

In this activity, we focused on analysing the possible stages of conflict that could affect the execution of the recovery process for the CIS BMS. It has been recognised and accepted that the priorities for information change in response to changes in the level of hostility and type of conflict that exists. Today's focus is on digitisation of the battlefield. The force that is active in the digitised environment has certain characteristics, as summarised below:

- Forces are sophisticated, professional and highly trained for short, high-tempo operations;
- The battle is likely to be non-linear, dispersed, and simultaneous;
- Battles will be of high intensity fought, if possible, away from the urban environment and population;
- Expeditionary deployment is likely; and
- Advanced technologies will be available – requiring an large investment in platforms and systems.

The characteristics of this digitised military, as well as the type of battle being waged and the spectrum of the conflict must all be carefully considered when developing the overall recovery strategy. We derived the stages of conflict to be: peacetime, pre-conflict, conflict, and post-conflict. These stages are applicable regardless of the type of conflict (e.g., civil, military). Because the scope of this research programme is on the military Command and Control (C2) environment, we focused on the stages of a conflict versus the types of conflict. As contained in the next section, we used these stages of conflict to evaluate potential changes in priority for recovery of key business processes.

5.2 *Priorities for the Key Business Processes*

The key business functions of the CIS span and map to the entire command structure. This is an important point to keep in mind when developing an overall recovery scheme. The first step in defining the recovery scheme for the key business processes was to have a hierarchy of most important (critical) to least important processes. In building this hierarchy, there were several considerations that had to be made including:

- Who you are
- What is your mission
- Where you are
- Who is under your command

The responses to the above queries impact the design of the hierarchy and will ultimately impact the recovery scheme itself. Also impacting the design of the hierarchy are two other issues: the type of conflict that you are engaged in (e.g., civil action, regional conflict, coalition warfare) and what stage of the conflict you are in (e.g., Peacetime, Pre-Conflict, Conflict, Post-Conflict). Figure 3 is an initial draft of the priorities for the key business processes. The list of key business processes was provided by DERA and originated under ARP19k. This example shows how the priorities can and will change during the different phases of conflict.

Peacetime	Pre-Conflict	Conflict/War	Post-Conflict
↓	↑ (p)	Combat Operations	↓
↔	↑ (p)	Joint Operational Picture	↔
↑	↑ (p)	Provide Intelligence	↔
↔	↑ (p)	Logistics/Movement	↓
↑	↑ (p)	Targeting	↓
↔	↔	Communications and Information Systems	↑
↔	↔	CCIRM	↑
↔	↑ (p)	Logistics/Supply	↑
↔	↔	Logistics/Medical Spt	↑
↔	↔	Personnel Spt	↑
↑	↔	Logistics/Maintenance	↑
↑	↑ (p)	Plans	↔
↑	↔	POLAD	↑
↔	↔	Operational Analysis	↑

(p) denotes the Planning aspects of these functions

Figure 3. Mapping of Key Business Processes to Changing Priorities

This table uses the current, key business process breakdown specified in the DERA-provided scenario for the conflict/war stage of a regional battle. We expanded this scenario to the other phases of a regional battle: Peacetime, Pre-Conflict, and Post-Conflict. We then analysed the processes in terms of these phases and developed a relative hierarchy for recovery. It is important to explain what we mean by “relative hierarchy.” Within the scope of this programme and with the definitions of the key business processes in flux, we did not feel that it would be appropriate, useful, or ultimately accurate to define a fixed hierarchy. In other words, we made no attempt to number the processes from 1 to 14, and change the numbering scheme within each process. Trying to do that would be akin to trying to put together a puzzle using pieces from 14 different and unrelated puzzles.

Instead, we derived a relative weighting that used the baseline scenario for Conflict/War and indicated the possible change in hierarchy using up and down arrows. For example, let’s look at the Combat Operations business process. During the Conflict/War stage, Combat Operations is the most critical function. (How could a conflict be staged without it?) During Peacetime though, the relative importance of Combat Operations will decrease because it is not an activity that is critical. The importance of Combat Operations – in terms of its Planning aspect – will increase during the Pre-Conflict stage, when the command is gearing up for a war. And, after the war is over, the importance of Combat Operations will again decrease due to its lower criticality. This same analysis was performed on all 14 key business processes and the results shown in Figure 3.