Cyber Risk to Mission: Assessment Methodology

Dr. David S. Alberts
Senior Fellow
Institute for Defense Analyses
Alexandria, Virginia
UNITED STATES

dalberts@ida.org

ABSTRACT

Today, we enjoy the benefits of a cyber-enabled force. However, as we have adapted our organizations, approaches to command and control, task processes, and doctrine to leverage cyber and cyber-enabled capabilities to increase our warfighting advantage, we have become more and more dependent upon these capabilities. This advantage is at risk if we are unable to ensure that our cyber and cyber-enabled capabilities are there whenever and wherever we need them. Perhaps we are too dependent upon these capabilities. In fact, a significant loss of cyber and cyber-enabled capabilities may put us at a greater disadvantage than if we had not deployed our cyber capabilities in the first place. This is because, over the years, we have adapted and would find it difficult to function without these capabilities.

Given our dependence on cyber-related technologies and the extent to which they are embedded in our operational capabilities, a deeper understanding of the portion of the fight taking place in Cyberspace and how it can impact the operational ecosystem is required. An understanding of the cyber health and status of the networks and cyber-enabled platforms and systems upon which operations rely is just as important as an understanding of their physical attributes, limitations, and readiness.

1. CYBER RISK TO MISSION (CRM)

CRM is a mission-focused concept that focuses on understanding the consequences that result from adversely impacted cyber capabilities. It involves assessing the likelihood of events that can result in a loss of cyber and cyber-enabled capabilities and understanding the linkages between these cyber events and measures of mission effectiveness.

CRM is present whenever the cyber or cyber-enabled capabilities upon which a commander depends fail to match operational expectations, putting the mission at risk.

2. RISK

Risk is commonly defined as an exposure to an undesirable circumstance or outcome. Two factors determine the extent of the exposure: first, the likelihood of an event or set of events measured by the probability that these events will occur; second, the significance of the consequences that flows from the event(s). When assessing CRM, the consequences of interest are those that impact our ability to carry out the mission(s).

Although risk is usually associated with damage and loss, a failure to capitalize on an opportunity is also a risk as it results in an undesirable outcome or "opportunity loss."

Different combinations of likelihood and consequences pose different types of risk. Events of concern can be rare or common; consequences can range from having an insignificant impact on a mission to having a catastrophic impact, making it impossible to successfully carry out the mission. Understanding the various types of risk is important as these risks vary in the priority, the urgency, and in the manner with which they need to be addressed.

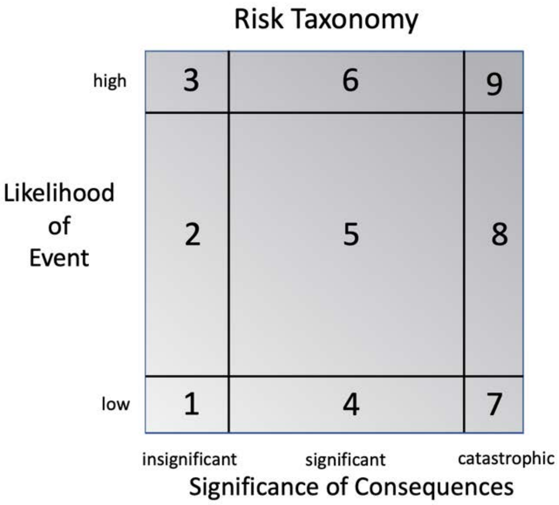

Figure 1: Risk Taxonomy

The Risk Taxonomy[1], depicted in Figure 1, considers both the likelihood of a loss[2] and its consequences to define nine Risk Types that serve to distinguish between and among the nature of, in this case, Cyber Risk to Mission. Likelihood of loss is assessed separately from the mission consequences because each provides valuable information that help develop an appropriate approach to manage risk. The boxes are sized differently to reflect the population distributions for likelihoods and consequences. Each of these distributions is assumed to have tails (extreme values are seen less frequently than "average" ones). This taxonomy can be used as a means to describe CRM and characterize a given approach, process, tools, or system capability to reduce risk as well as provide evidenced-based recommendations for managing CRM.

3. CYBER RISK

Cyber Risk to Mission is an "All Hazard" Risk; a shortfall in cyber and/or cyber-enabled capability can result from a variety of events or causes. These events include not only adversary cyberattacks but also include kinetic attacks, accidents, natural events, and/or system malfunctions.

Losses of cyber capability include adverse impacts to one or more of the following: availability, functionality, performance, assurance, confidentiality, integrity, security of, and/or our confidence in our cyber capabilities and the information they provide. Threats to the cyber capabilities upon which we depend may thus come from many sources and may manifest themselves throughout the competition continuum,[3] including periods in which events are characterized as being below the threshold of armed conflict. These threats include:

[1] Taken from Alberts, D. S. The Agility Advantage (2011)

[2] The likelihood of a loss considers both the likelihood of an event (hazard or attack) or events that could result in a loss and the likelihood that as a result of this event a loss occurs.

[3] Adverse cyber events often occur in situations that fall below the threshold of armed conflict. These events are important to understand and assess as they can create conditions that adversely impact preparations of missions, enabled subsequent attacks, and/or manifest themselves at a later time.

- Adversary actions
- Collateral damage from defending against real or imagined adversary actions
- Characteristics/complexities of cyber capabilities (hardware v. software/ software design and coding)
- Unanticipated behavior of systems, “intelligent” software, and decision aids
- Volatility of the cyber environment (rate of change)
- Collateral damage from cyberattacks on others
- Mistakes, accidents, poor cyber hygiene
- Critical infrastructure damage, degradation, disruption, denial, destruction

A CRM assessment needs, therefore, to consider a variety of events, each with a different potential to result in a loss of cyber or cyber-enabled capability. Some of these events could be part of an adversary "campaign plan” and be sequenced or orchestrated.

4. MANAGING CRM: REMEDIATION AND MITIGATION

Managing risk to mission requires “actions taken to remediate or mitigate risk or reconstitute capability in the event of loss or degradation”[1,2] CRM is best managed by a combination of remediation and mitigation. U.S. Department of Defense (DoD) definitions for remediation and mitigation are:

- Mitigation: Actions taken in response to a warning or after an incident occurs that are intended to lessen the potentially adverse effects on a given military operation or infrastructure.
- Remediation: Actions taken to correct known deficiencies and weaknesses once a vulnerability has been identified.

Another response would be to “Accept” the risk, if deemed appropriate. If the risk is accepted, that is, deemed not to require any immediate action, the situation should be appropriately monitored to continually assess the risk and address it, if and when appropriate.

Remediations and Mitigations of interest include not only technology-related solutions, but also those that include changes to organization, doctrine, command and control approach, processes, education, training, acquisition, and supply chain (as well as others).

5. CRM MEASURES

The bottom line of any Cyber Risk to Mission assessment is whether the risk is acceptable to mission commanders or appropriate decision makers. Many factors contribute to a determination of the threshold that separates acceptable from unacceptable risk. The CRM Assessment Methodology is therefore designed to provide decision makers with the information about the risk to mission necessary to make an informed choice.

6. CRM METRIC AND CYBER AGILITY

The concept of Cyber Agility,[4] “the capability that enables entities to succeed despite rapid, unanticipated loss of cyber or cyber-enabled capability which would otherwise threaten mission success,” can be used to develop

[4] The concept of Cyber Agility is based upon the publications of the DoD Command and Control Research Program.

a more complete understanding of the extent to which an organization is, was, or will be prepared to carry out its assigned tasks in the face of an unspecified set of cyber events and mission circumstances.

Cyber Agility is important to assess because, while understanding the CRM for a given mission and specific threats/hazards (scenario) is important and useful, it cannot be used as a substitute for an overall assessment of “cyber readiness,” since it represents only one possible future among many.

To develop an adequate understanding of CRM, one needs to consider more than one mission and more than a single threat-hazard scenario. This can be accomplished by considering both a “Threat-Hazard Space” that represents the set of possible hazards, threats, and mission circumstances and a “Mission Space” that represents the set of missions that could be undertaken. The ability to be successful over a Threat-Hazard Space is referred to as Cyber Threat-Hazard Agility while the ability to be successful in many regions of the Mission Space is referred to a Cyber Mission Agility.

Cyber Agility (a function of both Cyber Threat-Hazard and Mission Agility) serves as the overall metric for CRM. This metric can be employed for a single mission or for a collection of missions subjected to a set of possible threats/hazards and circumstance. Note that in the definition of Cyber Agility, the loss is unspecified. This is because it is the mission consequences that are of utmost interest.

7. CRM MODEL

Figure 2 provides a conceptual model of CRM upon which the assessment approach is based. This CRM model provides a depiction of the CRM for: 1) a single mission and set of circumstances (specification of a specific set of threats-hazards); 2) a single mission for a Threat-Hazard Space; and 3) CRM for a Mission Space.

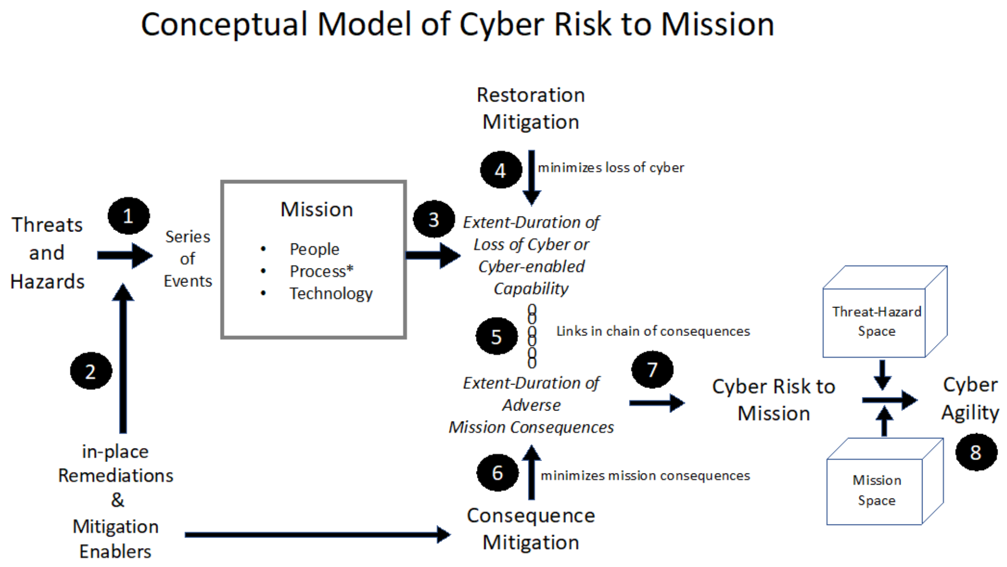
Conceptual Model of Cyber Risk to Mission
Restoration Mitigation
4
minimizes loss of cyber
Threats and Hazards
1
Series of Events
Mission
• People
• Process*
• Technology
3
Extent-Duration of Loss of Cyber or Cyber-enabled Capability
5
Links in chain of consequences
Threat-Hazard Space
2
Extent-Duration of Adverse Mission Consequences
7
Cyber Risk to Mission
Cyber Agility
8
Mission Space
in-place Remediations & Mitigation Enablers
6
minimizes mission consequences
Consequence Mitigation

*the C2 Approach defines the processes an organization employs to accomplish a given task or mission

Figure 2: Conceptual Model of Cyber Risk to Mission

Understanding the CRM associated with a single mission for a set of specified threats-hazards is the basic building block in constructing a CRM Assessment involving a space of possible or credible threats, and for assessments involving multiple missions.

As stated previously, Risk is a function of both likelihood and consequence. The likelihood of threats and risks can be incorporated into the Threat-Hazard Space while the likelihood and relative priorities of missions can be factored into the Mission Space.

Each component of the CRM Conceptual Model is discussed below.

1 Threats and Hazards

CRM stems from a set of possible threats and hazards that are capable of triggering events that result in a loss of cyber or cyber-enabled capability. Case studies and many wargames consider a specific set of threats-hazards while CRM assessments that employ models can consider a wide variety of potential threats-hazards (see discussion of Threat-Hazard Space below).

2 Remediations

The ability to prevent specific hazards/threat scenarios from causing damage sufficient to result in loss of cyber capability is determined by remediations that are “in place”; that is, those that are designed and built into technology, processes, C2 Approaches, and people.

Mitigation enablers can also be put in place; these either make it possible to take actions that can arrest or lessen the consequences associated with a loss of capability or facilitate the selection and/or taking of these actions. The mitigations themselves come into play after a loss is sustained. Mitigation enablers play an important role in determining the mitigation actions that are taken.

3 Extent and Duration of the Loss of Cyber and/or Cyber Enabled Capability

Figure 3 illustrates[5] the actual level of cyber and/or cyber-enabled capabilities over time (the solid black line), considering just one event and no restoration or consequence mitigations. This line constitutes the Zero Baseline to be compared with a line that represents the outcomes associated with proposed remediations and mitigations. Figure 3 takes into consideration remediations that were in place before the event occurred. The dotted line depicts the level of capability that would have been provided had not the cyber event taken place, or if it resulted in no significant damage.

[5] The shape of this line will be determined by the nature of the event and in-place remediations. It may, for example, “dive” rapidly or it may degrade slowly.

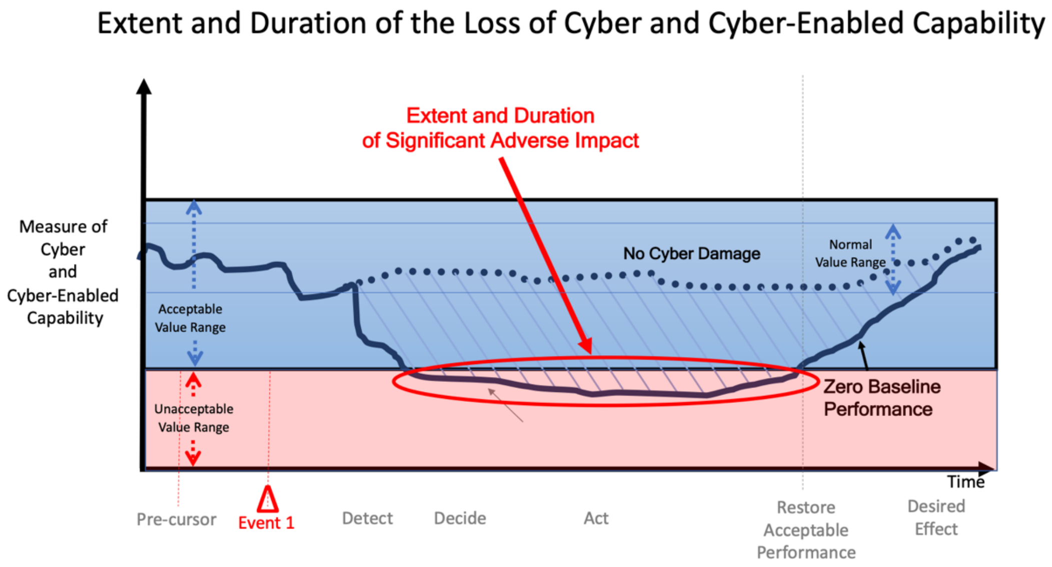

Figure 3: Loss of Capability over Time Given One Adverse Event

The extent and duration of the loss of cyber or cyber-enabled capability resulting from an event is thus a function of both remediations that are in place before cyber events occur and the organization, technology, processes, and people who work together to accomplish a given mission.

A value on the Y Axis denotes the threshold below which the cyber and cyber-enabled capabilities provided are no longer "acceptable." This value is derived from requirements or analysis.

Figure 3 also includes a number of points along the X Axis of interest. Starting with the left side of the graphic is the time(s) that one or more pre-cursor events are detected. Pre-cursor events are indicators that could provide warning of an adverse event that is likely to occur and provides an opportunity to take actions to prevent or lessen the extent and duration of the damage, or prepare to mitigate the consequences should the event take place. Moving to the right, the next time of interest is the time of the event itself, followed by the time it was detected. After detection, it takes time to gather necessary information, decide what to do, and act. After action is initiated, it takes time to have the desired effect of restoration of some or all of the loss capability. This time-line is important because it provides insight into the responsiveness of the restoration process that can be used to make the process more timely.

It is reasonable to expect that multiple "events" will occur; some of which will have been orchestrated to create enhanced adverse effects and present difficult challenges for designing, implementing, and prioritizing remediations and mitigations.

Figure 4 depicts a scenario that consists of two cyber events. Depending upon the nature of the events, the mitigation efforts necessary to effectively respond may impact each other. As with the previous figures, the graph is illustrative and the shape of the line in actual situations may take on many different shapes.

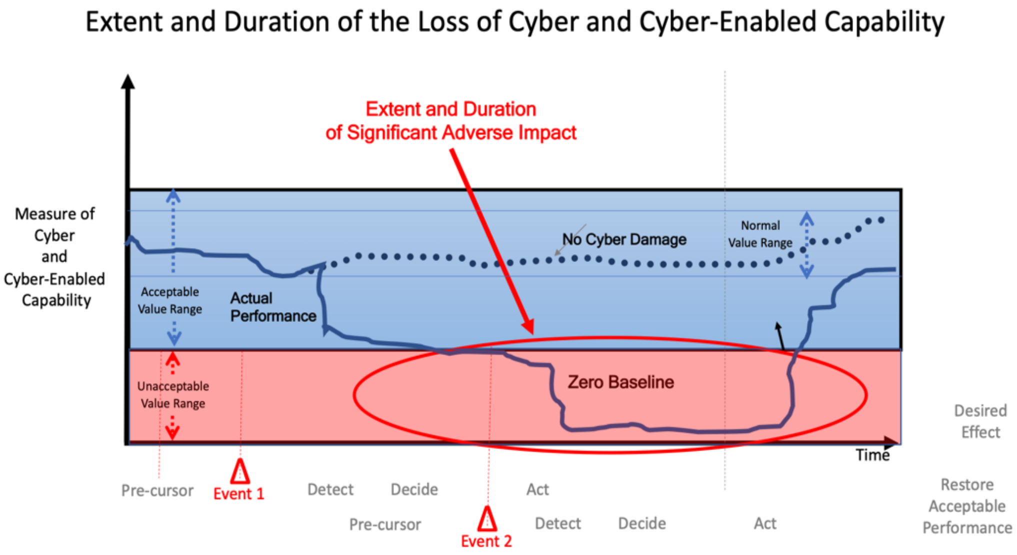

Figure 4: Loss of Cyber Capability Given Two Adverse Events

Figure 5 inserts a third "performance zone" defining performance levels that, while acceptable, constitute reasons for concern. Performance that falls in this zone of "heightened risk" indicates that, for the specified set of threats-hazards under consideration,[6] there is a need to prepare to ensure that mitigations are primed in case performance continues to degrade.

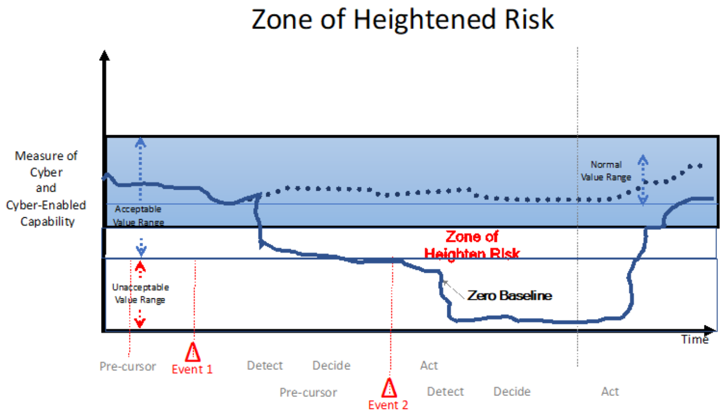

Figure 5: Zone of Heightened Risk

4 Restoration Mitigation

Once a loss has been sustained, efforts at restoration mitigation come into play to reduce both the extent and duration of the loss of cyber and cyber-enabled capability. The employment of Cyber Protection Teams or local defenders or system administrators are examples of a restoration capability.

[6] Falling into this zone does not affect the likelihood of events, but it does impact the severity of the consequences should the events occur, and hence, the CRM is increased.

The success at restoration efforts can be visualized in Figure 6. As with the previous graphs, the shape of the line that reflects the effectiveness of mitigation efforts will be a function of the situation.

In this illustration, the two events create the need for two separate restoration mitigation efforts, each of which does not bear fruit until the events are detected and a response is conceived and implemented. The first of these restoration efforts manages to increase performance from unacceptable to acceptable but it remains in the zone of heightened risk until the second mitigation is put in place when performance is restored to an acceptable level.

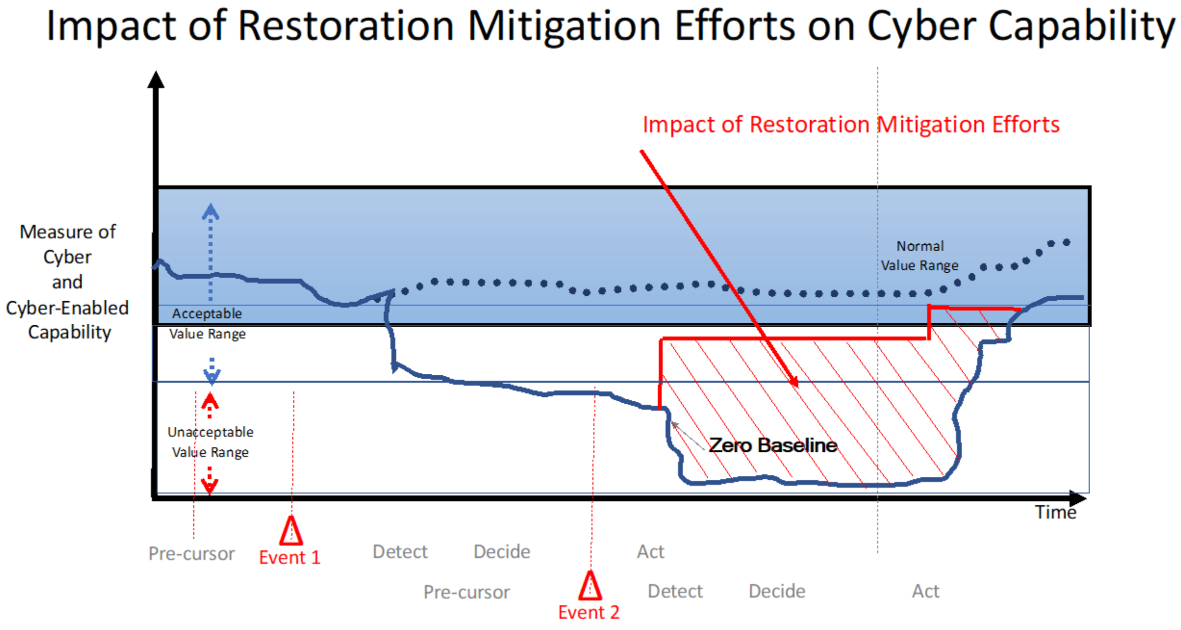
Impact of Restoration Mitigation Efforts on Cyber Capability
Impact of Restoration Mitigation Efforts
Measure of Cyber and Cyber-Enabled Capability
Normal Value Range
Acceptable Value Range
Unacceptable Value Range
Zero Baseline
Time
Pre-cursor
Event 1
Detect
Decide
Act
Pre-cursor
Event 2
Detect
Decide
Act

Figure 6: Impact of Restoration Mitigation Efforts

The net result of these restoration mitigation efforts is to significantly reduce the time period when performance is unacceptable. There remains, however, a significant amount of time where performance is at a level that is associated with a heightened risk.

Figure 7 provides a simplified three-value depiction of the extent of the loss of cyber capability with and without Restoration Mitigation. In this case, for the Zero Baseline, there was a long period of time where performance was unacceptable, bracketed by a short period of heightened risk. This resulted in a severity of consequences rating of "red". As a result of the restoration mitigation efforts, the time spent in the unacceptable zone was greatly reduced and performance remained in the zone of heightened risk for a long period of time. This is reflected in Figure 7 with a severity of consequence rating of orange." As noted earlier, when the situations with the same events are compared, the likelihood of these events are constant (an apples to apples comparison) and the comparison reflects relative CRM.

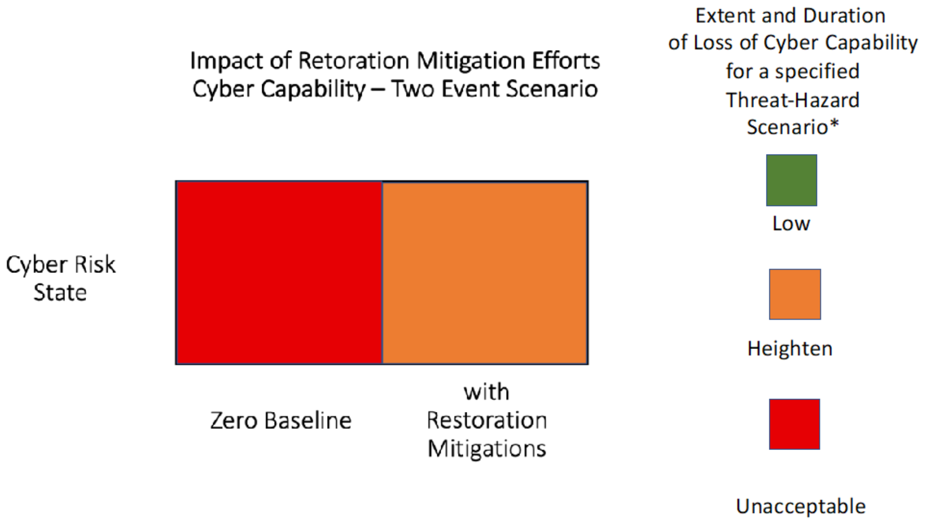
Impact of Retoration Mitigation Efforts
Cyber Capability – Two Event Scenario
Extent and Duration of Loss of Cyber Capability for a specified Threat-Hazard Scenario*
Low
Heighten
Unacceptable
Cyber Risk State
Zero Baseline
with Restoration Mitigations

*Both the Zero Baseline and the Situation with a set of Restoration Mitigations involve the same events and hence the Likelihood component of the risk calculus is unaffected

Figure 7: Simplified Depiction of the extent and duration of a Loss of Cyber Capability with/without Restoration Mitigation

5 Chains of Consequences

The effects of a loss of cyber and/or cyber-enabled capabilities (bottom two layers in Figure 8) spawns cascades or chains of consequences. Mitigation efforts are focused on arresting or managing the cascades of consequences that stem from a loss of cyber and cyber-enabled capability that may ultimately put a mission at risk.

Each of these chains impacts a different set of capabilities. The various threads that constitute the chain of consequences can be inter-dependent and hence constitute a network of dependencies where each link represents a dependency between a supporting capability and a supported capability, function, process or task. Figure 8 separates out specific capabilities into four layers. Each of these layers involve people, organizations, processes, technology and thus, each of these need to be represented and modeled so that impacts on and changes to any of these can be considered. The top layer contains mission processes and tasks. These are supported by cyber-enabled platforms and sensors and command and control.

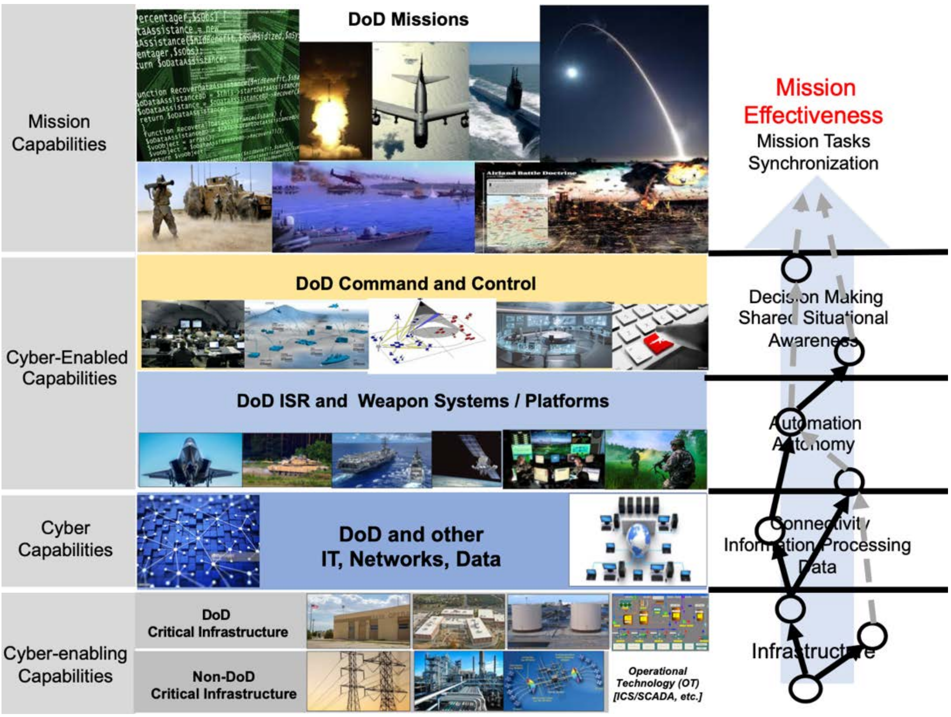

Figure 8: Capability Layer Model

For each "node" in this dependency network, there is an associated measure of performance and a graph that can be constructed that shows the level of performance over time (modeled after Figure 3). There is a cause-effect relationship between each link in the chain (and thus between adjacent graphs) that can be influenced or impacted by mitigation efforts. The focus of these capabilities and associated measures differ from layer to layer.

If fully effective and timely, restoration mitigation efforts can break one or more links in the chain of consequences before the impact is felt in the cyber-enabled or mission layers, thus arresting the adverse impact at that point protecting cyber-enabled capabilities from harm[7]. In doing so, no adverse impact would be felt in the mission layer. Short of this, these restoration mitigation efforts can shorten the duration and extent of the adverse impact as it cascades through the layers.

6 Consequence Mitigation

It is inevitable that there will be times when restoration mitigation efforts alone cannot stem the flow of consequences. The job of assuring the mission then falls to consequence mitigation.

Every node and link in the Dependency Network offers an opportunity to lesson or arrest a cascade of consequences. Efforts at Consequence Mitigation include a wide variety of efforts that can involve changes to processes, command and control arrangements, and alternative ways of accomplishing functions or tasks.

[7] When mitigations are effective they change the relationship between a dependency and its consequences. The 'original' dependency remains but the consequences change. One still has lost a capability but it no longer has the same adverse impact.

7 Cyber Risk to Mission

The final graph in the set depicts the extent and duration of the adverse impact of the loss of cyber capability on a measure of mission effectiveness (the Y axis) that takes into account the full set of restoration and mitigation efforts.

Both the measure of mission performance and the values that separate the ranges of performance into Consequence Severity levels are ideally determined by the mission owner. Whether unacceptable mission performance amounts to unacceptable risk depends upon both the extent and duration that the value of the mission performance metric falls below the acceptable threshold.

While this final graphic alone provides insight into the nature of the risk to the mission of interest under a specified set of circumstances, the additional information provided by the series of graphs (at various nodes of the dependency network) enhance this understanding.

Figure 9 provides, at a glance, the impact of a single, two-event threat-hazard scenario on both Cyber Risk and Mission Consequences for a single mission (derived from the final graph in the series of consequence graphs) with and without the Restoration Mitigation efforts of the CPT Force.

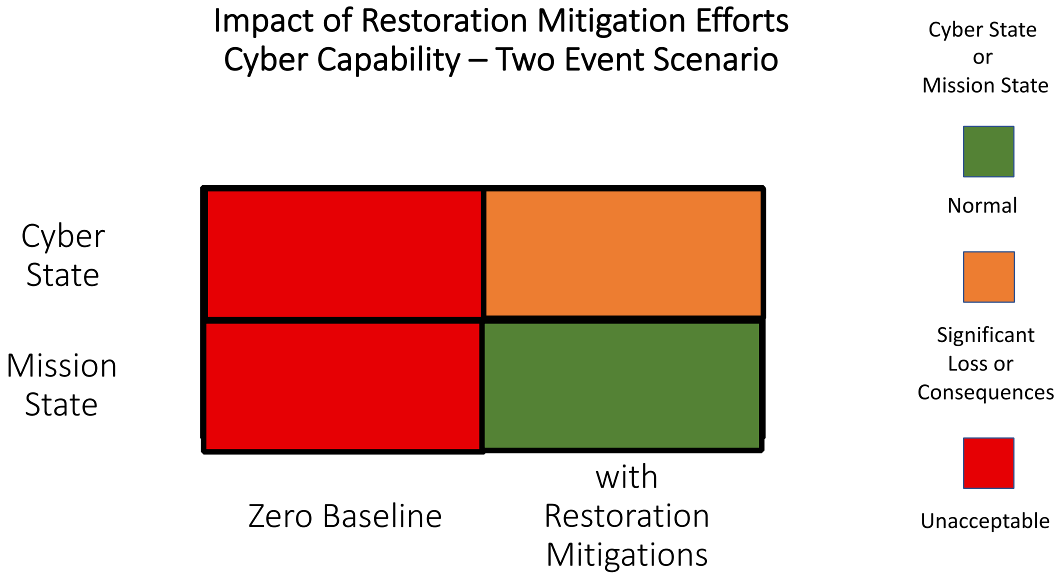
Impact of Restoration Mitigation Efforts
Cyber Capability – Two Event Scenario
Cyber State or Mission State
Normal
Significant Loss or Consequences
Unacceptable
Cyber State
Mission State
Zero Baseline
with Restoration Mitigations

*Both the Zero Baseline and the Situation with a set of Restoration Mitigations involve the same events and hence the Likelihood component of the risk calculus is unaffected

Figure 9: Cyber and Mission States with/without Restoration Mitigation

8 Cyber Agility

An ability to succeed in one mission for a given threat-hazard scenario equates to a situation (mission – threat-hazard scenario – circumstances) with low CRM.

Cyber Agility is the capability to remain successful despite a loss of cyber capability over both a Threat-Hazard Space (Cyber Threat Agility) and a Mission Space (Cyber Mission Agility).

Cyber Threat Agility

Cyber Threat Agility is the measure of risk for a given mission when anyone of a number of threat-hazard scenarios could occur.

A Threat-Hazard Space represents the nature and characteristics of the threats and hazards that are deemed to be credible. Regions of this space are representative of a set of threat-hazard scenarios that can result in similar adverse impacts to cyber capabilities and are thus associated with similar cascades of consequences. They, therefore, require a similar set for CRM management responses.

Specific hazards or threat scenarios can be mapped to region of the Threat-Hazard Space.

While each specific scenario is associated with a likelihood that would need to be accounted for in a scenario-based assessment of Cyber Risk, a Threat-Hazard Space based assessment only needs to consider the likelihood of challenges (regions of the space.) The likelihood of specific challenge would thus depend upon the likelihood of the most likely scenario that would present each particular challenge (region of the Threat-Hazard Space).

Cyber Mission Agility

Cyber Mission Agility is the measure of risk for a force that could undertake any number of missions.

One way to understand CRM when there are multiple missions of interest, is to look across the set of mission effectiveness extent-duration graphs for the specific set of missions and based upon the length of time a particular mission is in the various risk zones (low, heighten, or unacceptable), assign a CRM level to the mission.

Figure 10 illustrates the assignment of three risk levels to three illustrative graphs, given the same set of Threats-Hazards. Therefore, a change in the severity of the consequences will change the level of CRM.

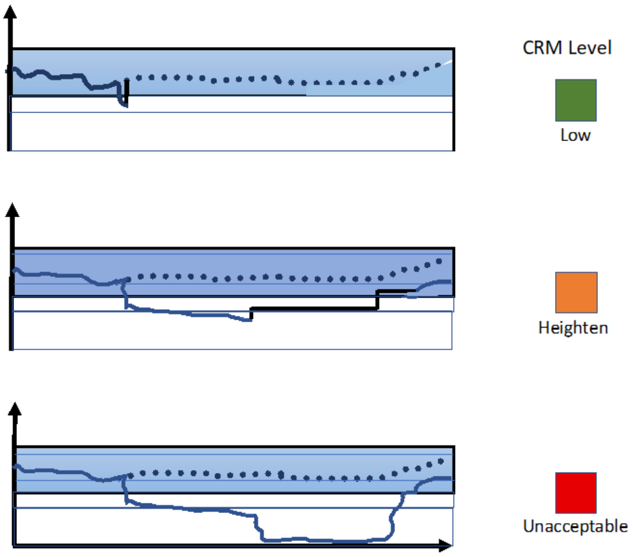

Figure 10: Assessing Risk Levels from Extent-Duration Graphs

Figure 11 depicts illustrative CRM Assessments for a set of missions. If the priorities and criticalities associated with these missions differ significantly, the missions can be grouped accordingly as depicted.

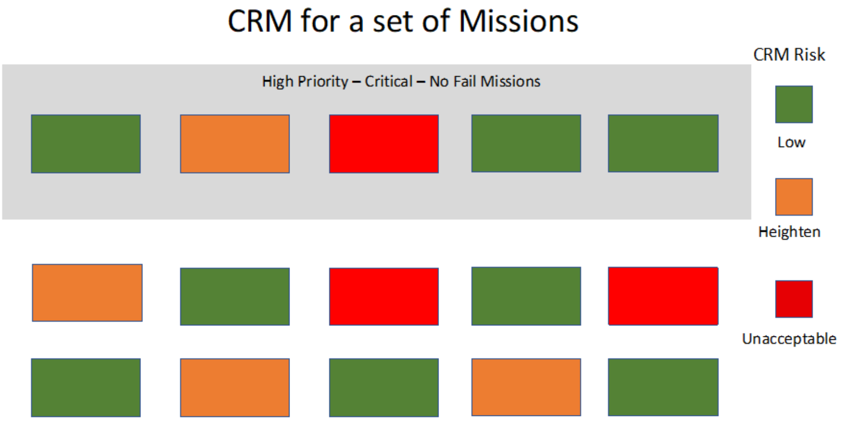
CRM for a set of Missions
High Priority – Critical – No Fail Missions
CRM Risk
Low
Heighten
Unacceptable

Figure 11: CRM Consequence Assessment for Missions of Interest

Figure 11 is one approach to visualizing the CRM for a set of missions of interest for a specified Threat-Hazard scenario. It could be used to focus attention on those missions that need to be addressed. It should be noted that the dependencies between and among these missions should be factored into rankings of the missions accordingly to criticality – priority.

The impact of a proposed change to remediation-mitigation capabilities can be depicted as in Figure 12.

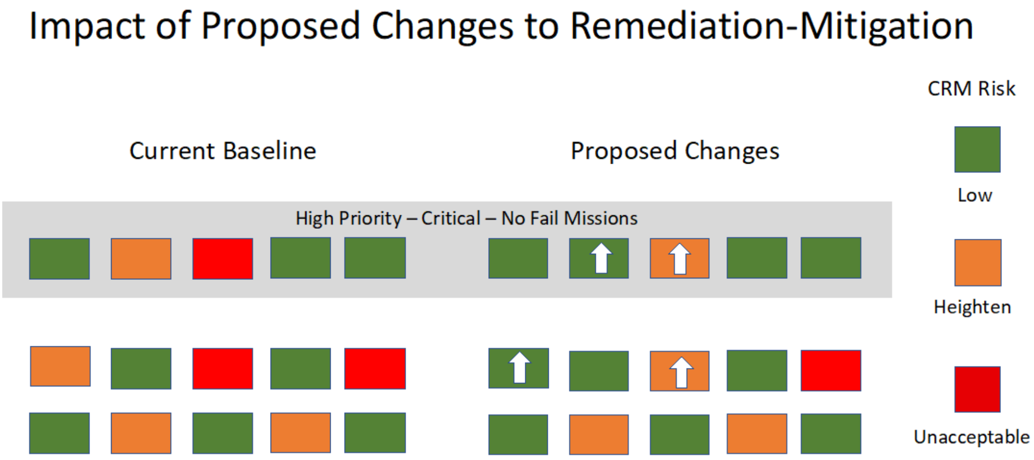
Impact of Proposed Changes to Remediation-Mitigation
CRM Risk
Current Baseline
Proposed Changes
High Priority – Critical – No Fail Missions
Low
Heighten
Unacceptable

Figure 12: Impact of Proposed Changes to Remediation-Mitigation

Construction of a Mission Space

An alternate approach to looking at a set of CRM assessments or a weighted sum that reflects their priorities is to create Mission Space and map the set of missions to this space. A Mission Space represents the nature and characteristics of missions. Regions of this space are representative of missions that have similar cyber dependencies and mission dynamics.

Figure 13 depicts a simple two-dimensional characterization of missions. One dimension is Mission Dynamics that reflects how quickly the situation can change and the size of time windows associated with opportunities to prepare, attack and/or defend. The other dimensions is Cyber Dependency that reflects how critical cyber enabled capabilities are to mission.

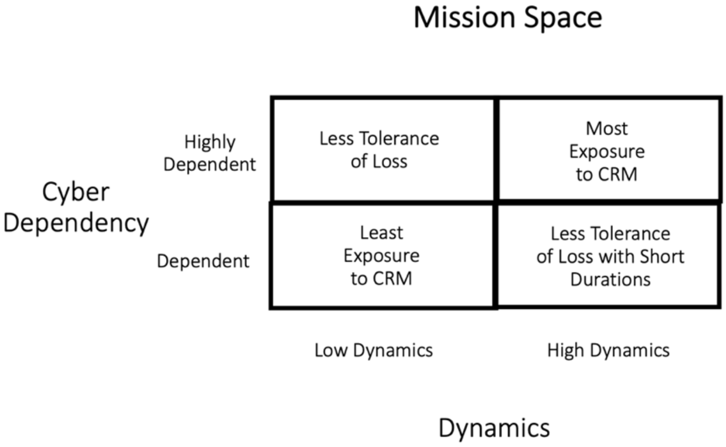

Figure 13: Mission Space

Missions that are both dynamic and are highly dependent on cyber capabilities for success will, of course, be most exposed to a loss of cyber capability even if the extent and duration is limited. Less dynamic missions will be able to tolerate losses of short duration, while missions that are not highly dependent on cyber capabilities (today all missions are, to one degree or another, relatively dependent) will be able to tolerate some loss of cyber capability. Missions that possess low dynamics and are not highly dependent on cyber capabilities are, of course, the least exposed.

Cyber Mission Agility is a function of the ability for a mission to remain successful in each of these quadrants of the Mission Space. Figure 14 depicts the severity of the consequences associated with each region of the Mission Space for a given CPT Force Alterative. Each region is colored according to the distribution of results from the DCM mapped to the region.

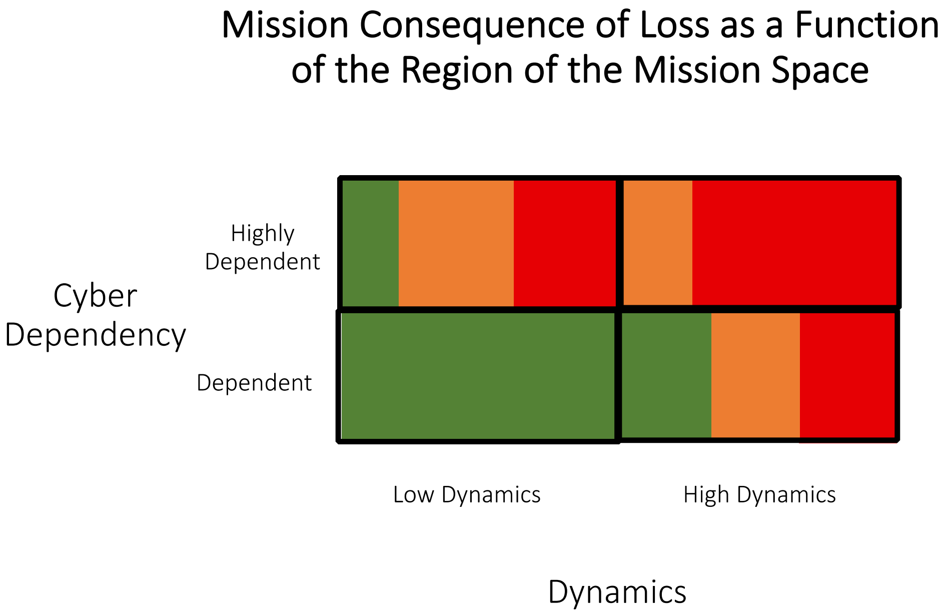
Mission Consequence of Loss as a Function of the Region of the Mission Space
Cyber Dependency
Highly Dependent
Dependent
Low Dynamics
High Dynamics
Dynamics

Figure 14: Severity of Mission Consequence as a Function of the Mission Space Region

8. SUMMARY

This paper presents a methodology and set of metrics that can be applied to a variety of Cyber Risk to Mission assessments. It expands the focus from looking at just losses of cyber capability to the consequences for missions. It enables a balanced approach to managing CRM as it provides an opportunity to understand the tradeoffs between remediation and mitigation.

As with any methodology, its value will depend upon an ability to populate it with credible data. Collecting appropriate data will require a “campaign” of analyses that include wargames, case studies, model development, and experiments that build upon each other.

Given the importance of cyber capabilities and the existence of a contested cyber environment, efforts to better understand CRM are urgently needed to assess the cyber readiness of our forces and to enable us to appropriately manage this risk.