MTR180449
MITRE TECHNICAL REPORT

MITRE

Cyber Resiliency Metrics and Scoring in Practice

Use Case Methodology and Examples

Deborah J. Bodeau
Richard D. Graubart
Rosalie McQuaid
John Woodill

September 2018

Dept. No.: T8A2
Project No.: 5118MC18-KA

Bedford, MA

Abstract

Cyber resiliency analysis and metrics are sensitive to a wide variety of assumptions about the operational, programmatic, architectural, and threat environments in which alternative solutions are identified and considered. The effective application of analytic methods, scoring, and metrics can be illustrated via use cases or notional worked examples. A cyber resiliency use case enables clear exposition of the problems, trade-offs, and metrics identified as part of a cyber resiliency analysis. This report presents the methodology – framework and process – for creating cyber resiliency use cases that was developed under the Measuring the Effectiveness of Cyber Resiliency research project. This work illustrates elements of the framework for a variety of possible use cases and presents the Vehicle Use Case in detail.

This page is intentionally blank. □

Table of Contents

List of Figures

List of Tables

1 Introduction □

Cyber resiliency is *the ability to anticipate, withstand, recover from, and adapt to adverse conditions, stresses, attacks, or compromises on cyber resources* [1]. It is increasingly an explicit concern at varying scopes or scales, ranging from components to critical infrastructure sectors, regions, and nations. Cyber resiliency for systems, missions, and programs is one aspect of trustworthiness to be addressed by systems security engineering [2]. In that context, systems engineers and architects seek ways to apply cyber resiliency concepts and to integrate resilience-enhancing technologies into architectures, designs, and operational systems [3] [4] [5] [6]. As they do so, they need to evaluate the relative effectiveness of architectural alternatives, as well as new technologies, products, or processes, for improving cyber resiliency and mission assurance. Similarly, program managers seek to determine whether investments in cyber resiliency will enable them to meet mission requirements more effectively.

A companion document [7] provides a general reference on cyber resiliency metrics and scoring methods. Cyber resiliency analysis and metrics are sensitive to a wide variety of assumptions about the operational, programmatic, and threat environments in which alternative solutions are identified and considered. Therefore, the effective application of analytic methods, scoring, and metrics can best be illustrated via use cases or notional worked examples.

A cyber resiliency use case is a notional worked example of how:

- Cyber resiliency concepts and constructs can be interpreted and applied to a representative situation;
- Cyber resiliency solutions can be defined for, or a specific solution or set of solutions can be applied to, that situation; and
- The relative effectiveness of alternative solutions can be compared in that situation.

Use cases illustrate how cyber resiliency can be applied in a variety of ways, depending on the situation (i.e., the mission, system architecture, threat model, risk management strategy, and programmatic constraints). For the Measuring the Effectiveness of Cyber Resiliency (MECR) project, a use case is also intended to illuminate how cyber resiliency metrics, measures of effectiveness, and scoring can be used to inform decisions.

A use case differs from a worked example in four key ways. First, the use case developer is typically a small team of cyber resiliency subject matter experts (SMEs) and SMEs in the technology or representative operational setting for the notional system. In a worked example, participants can include systems engineers, Program Office staff (if the example includes an acquisition program), and technologists, as well as cyber resiliency SMEs. Second, in a worked example, specific stakeholders are identified and their inputs are solicited. In a use case, the use case developer identifies and represents the concerns of different types of stakeholders. Third, in a worked example, a specific set of technologies and products are assembled in a system architecture. Depending on the stage in the system development lifecycle (SDLC), that architecture may be realized in an operational system or it may be represented via model-based engineering (MBE) or model-based systems engineering (MBSE). Finally, because a worked example involves a specific set of technologies, metrics and measures of effectiveness (MOEs) for alternative solutions can be evaluated in a model, laboratory, test environment, or operational setting.

A use case framework enables clear exposition of the problems, trade-offs, and metrics identified as part of a cyber resiliency analysis. This report presents two uses cases, one for a notional customer relationship management application and another for the acquisition of a vehicle fleet. It illustrates elements of the framework for a variety of possible use cases and presents the Vehicle Use Case in detail.

1.1 Methodology Overview

The Cyber Resiliency Use Case Methodology has two parts: a framework that identifies the types of information included in a use case, and a tailorable process for populating that framework. Figure 1 illustrates the five major components of the framework.

Figure 1. Use Case Framework

The components of a cyber resiliency use case are:

- **The programmatic and system use concept.** This component identifies the general type of system. The following types can be used in conjunction with the Cyber Resiliency Metrics Catalog [8] and the Cyber Resiliency Metric Template in Appendix C of [7]: cyber-physical system (CPS), enterprise information technology (EIT), large-scale processing environment (LSPE), and platform IT (PIT).

  The general type of system implies at a high level a concept of operations and an architecture. The system use concept also relates to the degree of aggregation or federation being considered, which has implications for governance, system administration, and information sharing between constituent systems in a system-of-systems. Finally, the system use concept includes a general characterization of the system user population and indicates whether cyber defenders are actively involved in monitoring the system and responding to indications and warnings (I&W) of adverse conditions or behaviors.

  The system use concept and/or the programmatic concept can identify a more specific representative of the system type to be examined in the use case. For example, a vehicle such as a car is a CPS, with multiple embedded control units (ECUs); however, significant differences exist between vehicles at different levels of automation.[1] Thus, a self-driving car would define a different use case from a car that provides no automation or driver assistance.

  This component also identifies the primary, secondary, and (if appropriate) supporting missions of the system, and the criticality of the system to those missions. Finally, this component identifies programmatic considerations for the acquisition, development, implementation or evolution of the system.

- **The architecture.** This component describes, at a high level, the system architecture and identifies interfaces with or dependencies on other systems. Note that "other systems" can include those constituting the system's development, test, or maintenance environment. Key technologies, technical standards, or products included (or expected to be included) in the system are identified. Locations, sub-systems or components, or layers in the architecture where cyber resiliency solutions could be applied are also identified and highlighted.

- **The threat model.** This component defines the characteristics and behaviors of adversaries whose attacks would undermine the system's ability to execute or support its missions, as well as the characteristics of relevant non-adversarial threats. Adversaries can include insiders as well as

[1] Six levels of vehicle automation are defined by the Society of Automotive Engineers (SAE). See https://www.nhtsa.gov/technology-innovation/automated-vehicles-safety.

individuals or groups located outside of the system's physical and logical security perimeter. Adversary goals are identified and translated into mission and cyber effects. Adversary behaviors – threat events, attack scenarios, or tactics, techniques, and procedures (TTPs) – are identified.

While a variety of threat scenarios can be defined for a given system, a use case can be restricted to a single scenario for expositive clarity.

- **Alternative solutions.** This component identifies the solution or set of alternative solutions considered in the use case. A cyber resiliency solution is a configuration of existing system resources or a set of technologies or products and supporting operational practices that solves the problem of reducing mission risk due to adverse conditions, stresses, attacks, or compromises on cyber resources. In the context of a cyber resiliency use case, alternative solutions are suggested by the motivating attack scenario(s), in conjunction with the system architecture and the identified constraints.

  A solution can be a single technology or architectural decision, or it can include multiple technologies, architectural decisions, and changes in operational or maintenance practices.

- **Scores, metrics, and MOEs.** This component identifies the scoring system or other method by which the overall assessment of the system is made, as a baseline and for alternative solutions. This component also identifies metrics and MOEs that could be or actually are evaluated to provide evidence to confirm (or disconfirm) the assessments of the alternative solutions.

  The Situated Scoring Methodology for Cyber Resiliency (SSM-CR, [7]) and the Cyber Resiliency Metrics Catalog [8] are used in this report.

These five components and the steps for developing them are described in more detail in Section 2.

Figure 2 illustrates the general process for constructing a cyber resiliency (CR) use case. The process consists of five steps, each of which can include multiple tasks. These are described in more detail in Section 2. The fifth step – identifying and evaluating relevant metrics and MOEs – is a key aspect for the MECR project. However, some use cases may only identify a representative set of metrics, for purposes of illustration.

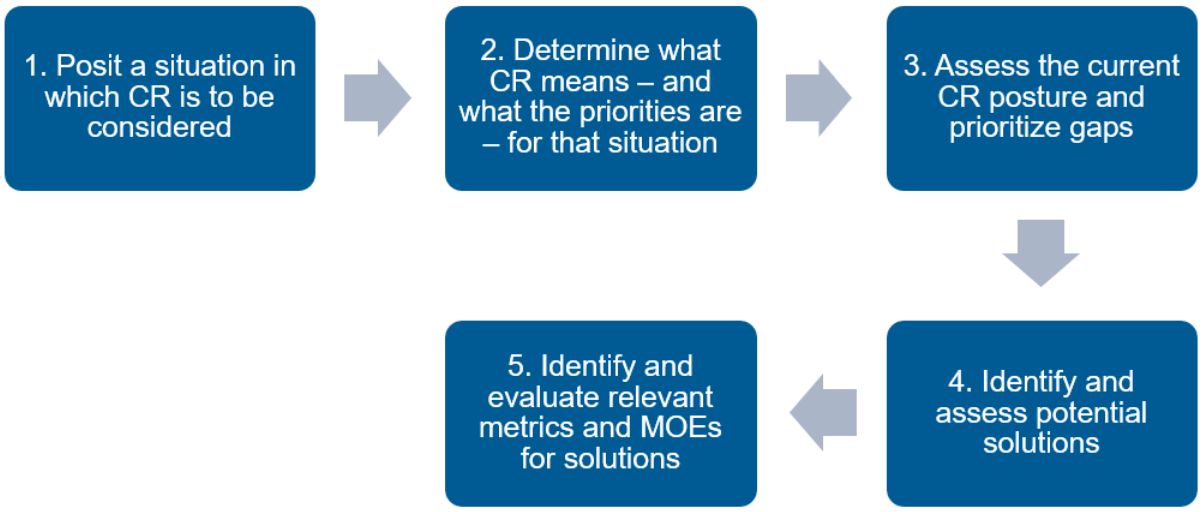

Figure 2. Use Case Development Process

1.2 Overview of This Document

Section 2 describes the use case development process in more detail, identifying tasks within the five major steps shown in Figure 2. Section 2 also describes the components of the use case framework,

showing how the steps of the use case development process populate those components. Section 3 presents a use case for a Customer Relationship Management (CRM) application, while Section 4 presents a vehicle use case. Details of the scoring for the vehicle use case are presented in Appendix A.

2 Use Case Development Process

The use case development process is shown in more detail in Figure 3. For the first step – positing a situation – all the more specific tasks are performed. For the remaining four steps, the more specific tasks selected will depend on the situation (as described in the first step) and on the level of detail sought for the use case.

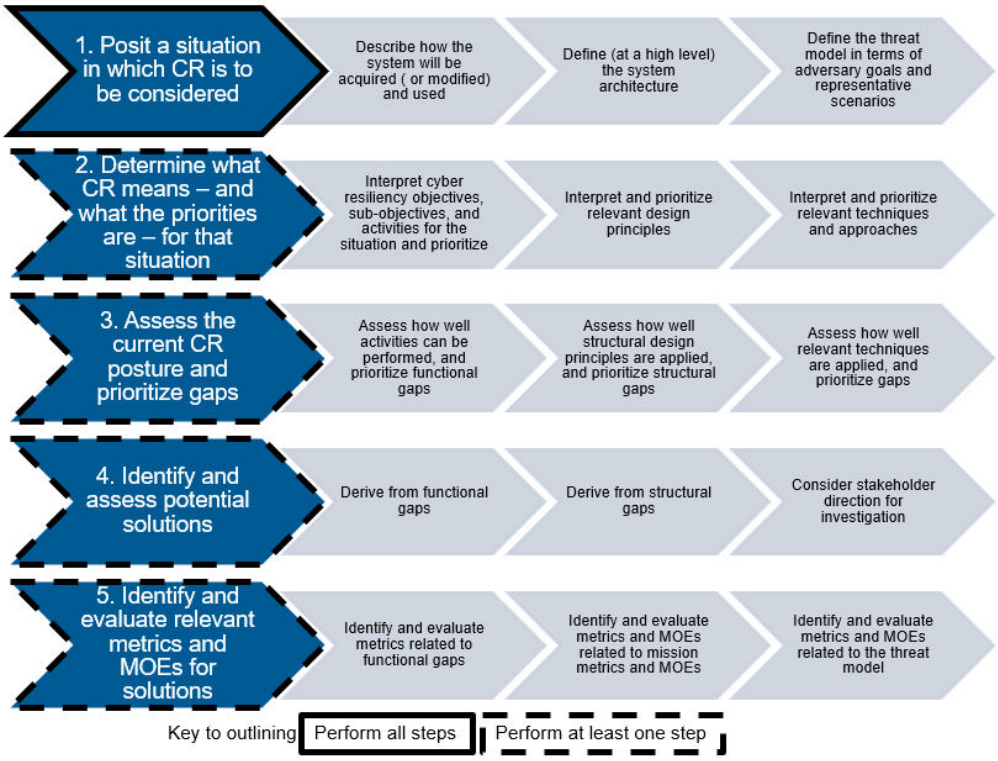

Figure 3. Use Case Development Process

Steps 2, 3, and 4 of the process use the Situated Scoring Methodology for Cyber Resiliency (SSM-CR) described in [7]. If another scoring system is preferred, it can be substituted. One of the tasks under step 5 of the process takes advantage of the Cyber Resiliency Metrics Catalog [8]; the tasks use metrics and MOEs related to mission or threat as described in Appendix A of [7].

2.1 Situate the Problem

The essential problem, in a cyber resiliency use case, is whether the cyber resiliency properties and behaviors of a system are sufficient for that system[2] to meet its mission assurance requirements. As illustrated in Figure 4, several questions must be addressed from the standpoint of the system's Program Manager or the Mission Owner of the mission the system supports, and the standpoint of the systems engineer or architect. Other questions also arise, related to how well the system meets its cybersecurity requirements, and whether alternative cyber resiliency solutions improve cybersecurity and overall system resilience, reduce existing cybersecurity risks, or can be traded off against cybersecurity requirements in light of mission assurance improvements. Such questions can be folded into discussions of the last two questions in Figure 4.

[2] "System" is construed broadly here, to include a system-of-systems (SoS), a constituent of an SoS, or an application or workflow which is separately acquired, managed, and maintained.

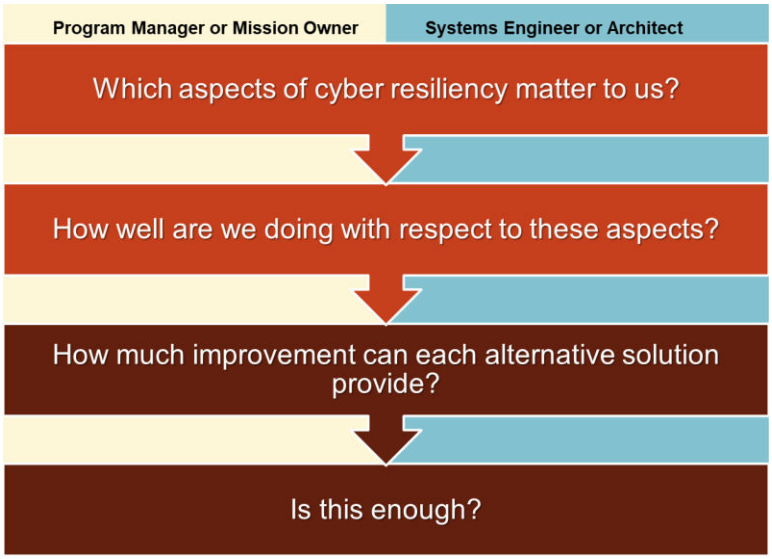

Figure 4. Key Questions in a Cyber Resiliency Use Case

These questions and the overall cyber resiliency problem are inherently situated in a context that includes the programmatic and system use concept, the system architecture, the underlying threat model, and existing capabilities which support system resilience, cybersecurity, or cyber resiliency. The first step in constructing a cyber resiliency use case involves defining this context.

2.1.1 Programmatic and System Use Concept

As indicated in Figure 1, the programmatic and system use concept is the first component of the use case framework. The developer of a cyber resiliency use case will:

- Define the system-of-interest.
  - Determine whether the system-of-interest is a system-of-systems, a constituent system within a system-of-systems, or a sub-system of a system.
  - Identify the general type of system – e.g., CPS (device, system, or system-of-systems), EIT, LSPE, PIT.
    - Multiple levels of aggregation have been defined for CPS: a device, a system, or a system-of-systems [9]. For example, a smart meter is an example of a CPS device; a vehicle is an example of a CPS; the Smart Grid is an example of a system-of-systems CPS. If the system is CPS, identify the level of aggregation.
    - Platform IT, as defined by the Department of Defense (DoD) [10], is typically a system-of-systems which includes both CPS and EIT.
    - Some systems are federated (e.g., the Smart Grid). Federation typically restricts the set of metrics which can be defined and used, since different system owners may be unwilling or unable to share certain types or forms of information. If the system is federated, note that fact.
    - Some systems are designed to operate without a network connection, at least transiently.[3] The set of cyber resiliency solutions, and the metrics which can be

[3] For example, many vehicles are designed to operate without a network connection, even though such a connection can be part of maintenance. However, that does not guarantee that the vehicle has no network connections, via its entertainment or

used to assess system cyber resiliency or solution effectiveness, will be limited by whether the system is operating in stand-off mode.

- Identify the primary mission or missions the system supports, any secondary or supporting missions, and the criticality and required reliability with which the mission is to be achieved.
  - Describe the system in terms of its intended uses, which include not only its primary mission or missions, but also secondary or likely additional uses. Identify external interfaces – to networks, to other supporting infrastructures and services, and to end users – keeping in mind that these interfaces can vary, depending on whether the system is operating under normal, stressed, or maintenance conditions, or whether the system is being used for one of its secondary purposes. Use this information to identify the system's attack surfaces.
  - Describe the system's criticality to its missions, its end users, or the general public. Criticality is "an attribute assigned to an asset that reflects its relative importance or necessity in achieving or contributing to the achievement of stated goals" [2], and relates most strongly to the potential impacts of system malfunction, degraded or denied performance, or mis-performance to the missions it supports, human life or safety, national security, or economic security (e.g., as in the context of critical infrastructure [11]).
  - If possible, identify mission MOEs and MOPs. Cyber resiliency effectiveness metrics can sometimes repurpose mission MOEs/MOPs, can sometimes repurpose data collected to evaluate mission MOEs/MOPs, and (particularly for cyber resiliency metrics related to Withstand or Recover) can often be related to mission MOEs/MOPs.
- Identify assumptions about the technical, operational, and decision environments. These assumptions constrain the set of possible cyber resiliency solutions.
  - The technical environment includes constraints on technologies and architectural decisions. It is determined primarily by the type of system, but is also influenced by aspects of the programmatic context as described below.
  - The operational environment includes the characteristics of the user community (e.g., size, awareness of cyber concerns), and of administrators and maintainers (e.g., general expertise, cyber expertise). It also includes whether the system is monitored by a Security Operations Center (SOC) and if so, how expert SOC personnel are. Another key aspect of the operational environment is how autonomous the system is, and whether its degree of autonomy changes over the course of mission execution.
  - The decision environment includes how cyber defense and mission contingency decisions are made – for example, a priori, in real time, or post hoc.
- Describe the programmatic context.
  - Identify the system development life cycle (SDLC) stage of the system,
  - Determine whether the development or maintenance environments are to be treated as part of the system.
  - Determine the extent to which legacy technologies and interfaces with existing systems must be accommodated. Identify requirements for interoperability with other systems or

navigation systems. A network connection can be via wireless (Wi-Fi), other radio frequency, or wired / optical networking, and need not rely on Internet Protocol (IP).

applications, incorporation of or support for legacy technologies, and functional dependencies on other systems or system elements.

- Identify (to the extent possible) the aspects of the programmatic risk management strategy which constrain possible solutions. One aspect is the relative priority of such quality attributes as safety, security, reliability, maintainability, system resilience, and cyber resiliency. Another is the relative preference for operational vs. technical changes – is it preferable to change how the system is implemented, or how it is used?

2.1.2 Architecture

The purpose of examining the system architecture in a cyber resiliency use case is to inform the development of threat scenarios (see below), enable cyber resiliency solutions to be described in terms of their architectural placement, and provide insight into possible metric evaluation. More specifically, the developer of a cyber resiliency use case will:

- Develop a conceptual view of the system, to establish the scope or bounds of the system and any possible solutions. This conceptual description of the system depends on the system type, its technical and operational environments, and its programmatic context. It can include aspects of the development and maintenance environments as well as the operational system-of-interest. The conceptual view relates to governance and the system concept of operations (CONOPS).
- Develop a technical view of the system, to describe how it is put together. This view identifies the key components or constituent sub-systems, and their interfaces, dependencies, and information flows. These can be targets of attack, and their interfaces and dependencies define the system's internal attack surface. This view also identifies the architectural layers included in the system-of-interest, and the assumed or demonstrated properties of layers on which the system-of-interest depends. This view can identify locations (architectural layers, components, flows between components) where solutions could be introduced, taking the constraints arising from the identified assumptions into consideration. Finally, this view can identify locations where observations could be made or data collected to serve as input into MOEs.
- Develop an operational view of the architecture, to describe how the architecture is exercise. This view identifies workflows for mission operations, system management, and cyber defense. It enable component criticality to be identified as a function of the mission workflow. It can also identify the system's modes of operation (e.g., optimal, sub-optimal, minimum essential, degraded or safe mode, graceful shutdown).
- Identify existing security, system resilience, and cyber resiliency controls, techniques, approaches, and/or design principles.
- Identify external interfaces and dependencies. Quality properties (e.g., security, reliability) of systems with which the system-of-interest interfaces, and in particular of those on which it depends, are identified.

Depending on the type of system and the applicable cyber resiliency objectives, design principles, or techniques, it may be possible at this point to identify some potential cyber resiliency metrics from the catalog [8].

2.1.3 Threat Model

Cyber resiliency solutions are motivated by threats to the system and the information it handles, the missions it supports, its users (both individual and institutional), and to the larger ecosystem of which it is a part. While these threats can be due to a variety of sources (e.g., human error, natural disaster, failure of a supporting infrastructure or service), a cyber resiliency use case focuses on adversarial threats, as

captured in a small set of motivating scenarios. Threats are described in terms of characteristics and behaviors. Characteristics include goals and intended cyber effects – why an adversary would attack as in a scenario. Behaviors are described in terms of adversary tactics, techniques, and procedures (TTPs), and can be categorized using the categories of the National Security Agency Central Security Service (NSA/CSS) Technical Cyber Threat Framework (NTCTF, [12]) or the ATT&CK framework [13].

The developer of a cyber resiliency use case will:

- Identify the types of threats considered in programmatic or organizational risk framing. As noted above, in addition to adversarial threats these can include threats of human error, faults and failures, and natural disaster. A use case can identify scenarios in which adversaries can take advantage of the consequences of non-adversarial threat events.
- Identify the adversary's characteristics, constructing an adversary profile.
    - Identify the adversary's ultimate goals and intended cyber effects. (See Table 9 of [14].)
    - Identify the timeframe over which the adversary operates.
    - Identify the adversary's persistence (or, alternately, how easily the adversary can be deterred, discouraged, or redirected to a different target).
    - Identify the adversary's concern for stealth. (See Table 10 of [14] for representative values for timeframe, persistence, and concern for stealth.)
    - Describe the adversary's targeting, which relates to the scope or scale of the effects the adversary intends to achieve. (See Table 11 of [15].)
- Identify the representative attack scenarios of concern, describing each scenario with a phrase or a sentence. A set of general attack scenarios identified in [15] [14] can serve as a starting point. The attack scenarios of concern in the cyber resiliency use case should be clearly related to the system's mission. Note that a use case can focus on a single attack scenario, or can consider a set of scenarios. The purpose of identifying the attack scenarios is not to be exhaustive – the use case is not a full-fledged risk analysis – but rather to serve as the basis for assessment and metric identification.
- Describe the representative attack scenarios, identifying stages in the attack (e.g., administer, engage, persist, cause effect, and maintain ongoing presence [12]) and the system elements compromised in each stage.
- Identify common elements across the attack scenarios (e.g., recurring adversary TTPs as defined by ATT&CK or [12]), as a starting point for identifying potential alternative solutions.

A use case can also include representative threat scenarios related to non-adversarial threat sources. For these,

- Identify the scope or scale of effects, duration or timeframe, and types of assets affected. See Section 5.1.2 of [15].
- If possible, provide a reference to a publicly available description of a similar scenario to serve as an anchoring example.

For purposes of illustration, a use case can focus on a single threat scenario.

2.2 Interpret and Prioritize Cyber Resiliency Constructs

The intent of this step is to answer the first question in Figure 4: *Which aspects of cyber resiliency matter to the stakeholders in the system, mission, or organization?* To develop that answer, the developer of the

use case must ensure that cyber resiliency concepts and constructs are meaningful in the context defined in the first step. The developer of the use case will do one or more of the following:

- Restate and prioritize cyber resiliency objectives, sub-objectives, and activities. These constructs are restated in terms meaningful to the architecture and the system use concept. They are prioritized based on programmatic considerations and mission concerns. Note that responsibility for some activities may be allocated to system elements outside the scope of what can be affected by decisions in the use case. The prioritization is an input to SSM-CR; for not-applicable objectives, sub-objectives and activities can be ignored in subsequent steps.
- Determine the potential applicability of cyber resiliency design principles, techniques, and implementation approaches. This involves considering organizational and programmatic risk management strategies to determine which strategic design principles may apply. It also involves considering the architecture, system use concept, and threat environment, to identify the relevance of structural design principles to this situation. Relevant structural design principles are restated in situation-specific terms (e.g., in terms of the technologies that are part of the system).
- Determine the potential applicability of cyber resiliency techniques and implementation approaches. This involves considering the architecture, system use concept, and threat environment. The relevance of techniques and approaches to this situation is described and assessed. Relevant techniques and approaches are restated and described in terms of architectural elements – e.g., allocating an implementation approach to a specific system element.

Prioritization is a key input to the baseline assessment and to the identification of potentially relevant metrics.

2.3 Perform Baseline Assessment and Prioritize Gaps

The intent of this step is to answer the second question in Figure 4: *How well is the system doing – how well does it meet stakeholder needs and address stakeholder concerns – with respect to the aspects of cyber resiliency that matter to stakeholders?* The developer of the use case will do one or more of the following:

- Assess how well relevant cyber resiliency activities, as restated and prioritized in the previous step, are or can be performed. The assessment can use the Situated Scoring Methodology for Cyber Resiliency (SSM-CR). In SSM-CR, SMEs assess the level of performance (on a scale from 0-5) for relevant activities; these assessments are combined ("rolled up") to produce assessments for relevant sub-objectives and objectives (as restated), and to provide an overall assessment of cyber resiliency for the system. See Figure 5 below and Appendix D.2 of [7] for more details.
- Assess how well the relevant Cyber Resiliency Design Principles (CRDP) have been applied. See Appendix D.3 of [7] for more details.
- Assess how well the relevant cyber resiliency techniques and approaches are applied. See Appendix D.3 of [7] for more details.

- **Assess the relative priority of each objective on a scale of 0-5. For each relevant (non-zero priority) objective, assess the relative priority of each sub-objective. For each relevant sub-objective, assess the relative priority of each activity. Note that relevant objectives, sub-objectives, and activities are interpreted or tailored to the situation.**
- **For each relevant activity, assess the performance level on a scale of 0-5.**
- **Roll up the assessments to produce a cyber resiliency score as follows:**
  - For each relevant sub-objective, performance level = $100 * (\sum_{activities} \text{Priority(activity)} * \text{Performance(activity)}) / (\sum_{activities} \text{Priority(activity)} * 5)$
    - If all activities have 0 priority, the denominator is set to 1; the result is 0.
    - This formula captures the percentage of the maximum priority-weighted performance achieved by the actual priority-weighted performance.
  - For each relevant objective, performance level = $100 * (\sum_{sub\text{-}objectives} \text{Priority(sub-objective)} * \text{Performance(sub-objective)}) / (\sum_{sub\text{-}objectives} \text{Priority(sub-objective)} * 100)$
    - If all sub-objectives have 0 priority, the denominator is set to 1; the result is 0.
    - This formula captures the percentage of the maximum priority-weighted degree of achievement achieved by the actual priority-weighted achievement.
  - For overall cyber resiliency, performance level = $100 * (\sum_{objectives} \text{Priority(objective)} * \text{Performance(objective)}) / (\sum_{objectives} \text{Priority(objective)} * 100)$
    - If all objectives have 0 priority, the denominator is set to 1; the result is 0.
    - This formula captures the percentage of the maximum priority-weighted degree of achievement achieved by the actual priority-weighted achievement.

The Cyber Resiliency Score is on a scale of 0-100. This is to be interpreted as a semi-quantitative value – useful for comparisons, but in no sense absolute or highly granular – since it is computed using semi-quantitative inputs. Thus, the range of 0-20 is Very Low; 21-40 is Low; 41-60 is Moderate; 61-80 is high; and 81-100 is Very High.

Figure 5. SSM-CR Scoring

The baseline assessment is situated in terms of the programmatic and system use concept, architecture, and threat model. For each of the most granular cyber resiliency constructs (i.e., activities, structural design principles, or approaches), the use case developer identifies gaps – differences between the ideal (e.g., a value of 5 or Very High) and the actual assessed value. Gaps in performance of activities can be prioritized, based on the relative importance of the sub-objectives and objectives they support, and the size of the gap. Gaps in the application of design principles can be prioritized based on the relevance of the design principle and on the size of the gap in how broadly and how well they have been applied. Similarly, gaps in the application of approaches can be prioritized based on the relevance of the approach (or the technique it implements) and on the size of the gap in how broadly and how well the approach has been applied.

In addition, this step can also include an assessment of identified gaps in cybersecurity controls, with respect to the severity to mission performance resulting from those gaps.

2.4 Identify Potential Solutions

A potential solution is a combination of products, technologies, or architectural decisions with operational processes or practices, which is expected to improve the system's cyber resiliency. The level of detail with which a potential solution is described depends on how specifically the situation was described in the first step. In particular, if the architecture and the operational environment were described in general terms, potential solutions will also be described at a high level.

The use case may specify a potential solution, or set of solutions, to be considered. If so, this step consists of characterizing that solution or set of solutions in terms of the cyber resiliency techniques and approaches (and/or the design principles) it applies.

Alternately, the developer of the use case can identify potential solutions by focusing on the threat scenario(s) in the threat model, analyzing the architecture to identify potential locations where a scenario could be interrupted or its effects reduced. The description of a potential solution can include identification of the gaps it is expected to address, as well as the potential locations at which it could be applied. Locations are identified using the technical view of the architecture (e.g., layers, components, interfaces). Potential solutions are analyzed for relevance (i.e., whether and how a proposed solution reduces risk from the motivating attack scenario(s)) and viability, in light of the identified constraints.

If more than one potential solution is determined to be relevant and viable, the developer of the use case analyzes these potential solutions for compatibility or interference. Solutions may be mutually supportive, independent, or incompatible.[4] A combination of mutually supportive or independent solutions can be defined as an additional solution.

The identification of potential solutions is supported by an analysis of which cyber resiliency design principles, techniques, and/or approaches could be applied. This supporting analysis can include restatements of relevant cyber resiliency objectives or design principles in terms of the type of system, its missions, and its architecture, if those restatements were not already made in the second step.

In addition, depending on the level of detail in the threat model, the description of a potential solution can include identification of the threat events or TTPs it is intended to affect (and what effects can be expected).

2.5 Assess and Measure Potential Solutions

The intent of this step is to answer the third question in Figure 4: *How much improvement could each alternative solution provide?* At a minimum, in this step the developer of the use case revisits the assessment from the third step for each potential solution, identifying changes to assessed values. In general, a potential solution is expected to improve or make no changes in assessed values. However, a potential solution can introduce the possibility of new threat scenarios; if this is the case, then some assessed values could decrease. In addition, the list of new threat scenarios introduced by a potential solution is a descriptive or nominal metric for its cost.

In addition, the developer of the use case can identify – and depending on the availability of evaluation environments, can evaluate – relevant metrics, without and with potential solutions. Metrics which indicate cyber resiliency properties can include measurements, observations, or values computed from these related to expected or observed system behavior or performance under adversity, mission performance under adversity, or system properties. Relevant metrics can be identified from the identified gaps, e.g., representative metrics identified for activities in Appendix B of [7] or representative metrics identified for design principles in [3]. These metrics must be tailored and specified for the posited situation

The developer of the use case can thus identify (and possibly evaluate) MOEs for potential solutions. MOEs can take the form of changes in values of indicator metrics, changes in mission MOPs or MOEs under adversity, or measurements of effects on adversary activities or threat scenarios.

Evaluation can be performed in any of a range of environments, depending on the use case. These can include operational systems, representative environments (e.g., cyber range), laboratories, emulations, or modeling environments. However, because most use cases are notional, evaluation of metrics and MOEs is typically either not performed, or performed in an emulation or modeling environment.

[4] See [4] [3] for more information.

3 CRM Use Case

The CRM use case, as summarized in Figure 6, involves the possible use of micro-segmentation to re-architect a customer relationship management system used by a large organization to manage claims processing. Claims are made by members of a large public; the CRM application manages the workflow to ensure that they are processed correctly, in a timely manner, and without fraudulent activity from outsiders or insiders. The CRM application runs on an established enterprise infrastructure; changes to that infrastructure are explicitly out of scope for this use case.

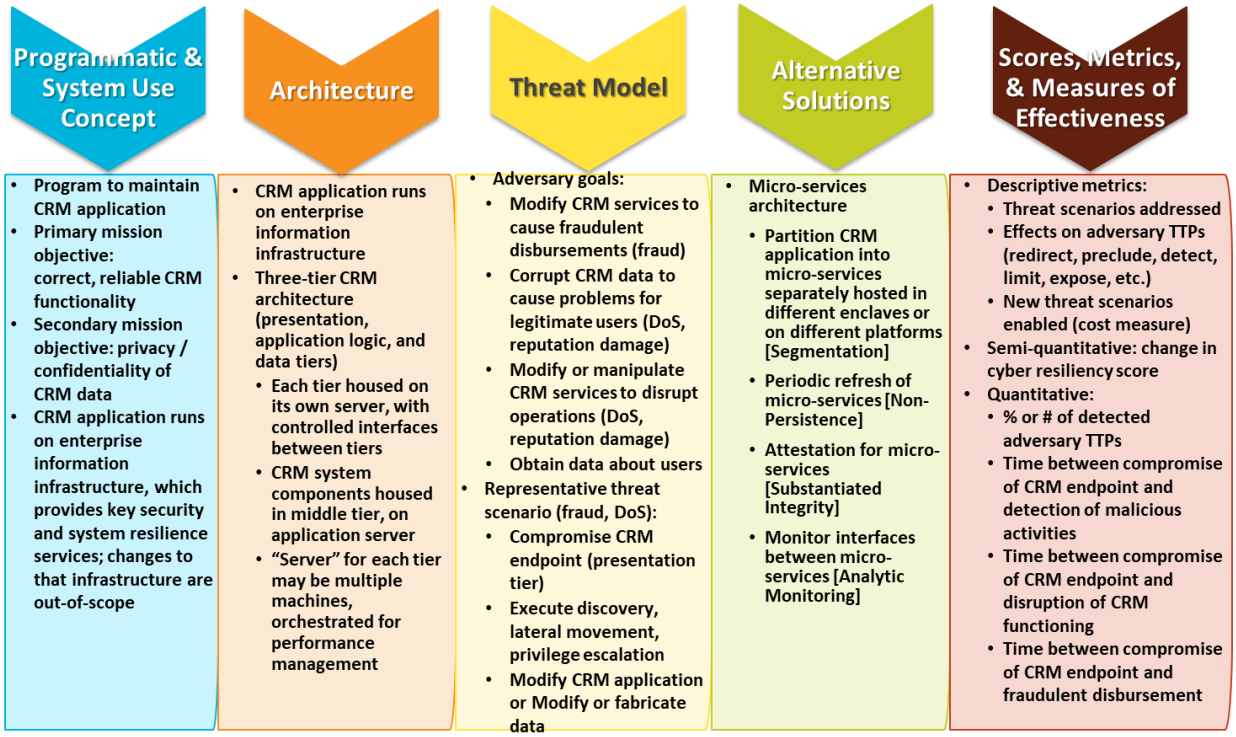

Figure 6. Summary of CRM Use Case

3.1 Programmatic and System Use Concept

The system-of-interest is the CRM application used for claims processing. Because changes to the enterprise infrastructure on which the CRM application runs are explicitly out of scope, and the focus is on whether micro-segmentation could improve cyber resiliency, the scope of this use case is quite narrow.

3.1.1 Mission

The primary mission of the CRM application is to manage cases created by and serving individuals in a large population, so that services and/or funds can be provided to those individuals in a timely way. The secondary missions are to maintain the privacy of served individuals, and to reduce the potential for fraud, waste, or abuse of organizational resources.

The criticality of the CRM application to its mission is very high; realistically, case management cannot be accomplished without the CRM application. In principle, the organization could fall back to a paper-based workflow; in practice, this would not scale to the served population, and would greatly increase the potential for fraud.

The system may be stressed in times of crisis (e.g., there may be a surge in claims following a natural disaster) or during specific predictable periods (e.g., immediately prior to filing deadlines).

3.1.2 Environmental Assumptions

Key technical assumptions are that:

- The CRM application runs on the enterprise infrastructure, which provides supporting security, networking, and performance management services.
- The CRM application must be compatible with the enterprise architecture. This includes interoperability with legacy applications (e.g., database management systems or DBMSs) and enterprise services, including identity and access management (IdAM), intrusion detection systems (IDSs), and insider threat monitoring.

The CRM application is currently operated as a single entity, with recognized risks as discussed with respect to the threat model described below.

Assumptions about the operational environment focus on the user population: a small set of administrators, focused on performance management; a moderate-sized population of enterprise staff (a.k.a. resources) responsible for claims management, varying in skills and roles, but with only basic cybersecurity training; a moderate-sized population of external partners, which can vary widely in cybersecurity capabilities; and a very large population of external users (potential claimants), who have no cybersecurity awareness and may well be using compromised devices to interact with organizational systems.

Cyber defense[5] is an enterprise service, provided by a low-capability security operations center (SOC).

3.1.3 Programmatic Constraints

The use case focuses on a planned upgrade to the CRM application, to improve its performance and maintain interoperability with enterprise services. That upgrade, like any effort related to the organization's information technology (IT), must respect the organization's risk management strategy, which strongly prioritizes reduction of fraud risks. Thus, the CRM Application Upgrade Program has a strong motivation to reduce risks from advanced cyber threats as well as insiders. However, the upgrade needs to avoid disrupting CRM service delivery or imposing new requirements on the enterprise infrastructure.

3.2 Architecture

The CRM application under consideration follows a three-tier architectural pattern, typical of many public-facing enterprise applications. As illustrated in Figure 7, the CRM application consists of three sub-systems:

[5] Active cyber defense is defined as the "synchronized, real-time capability to discover, detect, analyze, and mitigate threats and vulnerabilities." It is complemented by proactive cyber defense, "a continuous process to manage and harden devices and networks according to known best practices," and regenerative cyber defense, "The process for restoring capabilities after a successful, large scale cyberspace attack, ideally in a way that prevents future attacks of the same nature." [29] These different forms of cyber defense can support the cyber resiliency goals: proactive cyber defense can support Anticipate, active cyber defense can support Anticipate, Withstand, and Recover, and regenerative cyber defense can support Recover and Adapt. The extent to which an organization's cyber defense efforts actually support cyber resiliency depends on the maturity and capabilities of the organization's SOC, cyber security program, or information security program. An immature and poorly resourced program will focus on discovery and mitigation of vulnerabilities, and on incident or intrusion detection and response; while these activities are a necessary part of active cyber defense, they are far from sufficient to provide cyber resiliency.

- The presentation tier. This consists of a browser client to interact with end users (i.e., those making claims) and a custom client to interact with partners, running on an enterprise-provided User Interface Server.
- The application tier, which runs the application logic. The application consists of five major services:
  - Claim and case filing.
  - Workflow management.
  - Resource management. A resource, in this context, is an enterprise staff member responsible for some task or tasks in the claim or case management workflow.
  - Reporting.
  - Analytics.
- The data tier, which uses enterprise-provided storage and data management systems (e.g., a database management system or DBMS).

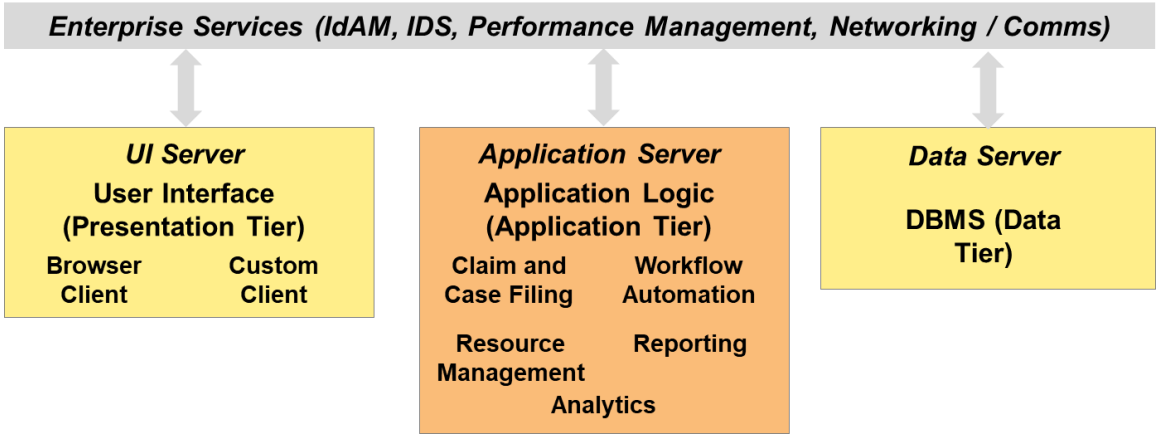

Figure 7. Three-Tier Architecture for Notional CRM Application

All communications between tiers use – and are monitored by – enterprise services. Currently, the application logic is managed as a single service. This severely limits the insight of enterprise services – for performance management as well as intrusion detection and insider threat monitoring – into the CRM application.

3.3 Threat Model

The focus of this use case is on adversarial threats, motivated by financial gain. Adversaries seek to defraud or steal money from the organization, by creating false claims, manipulating existing claims (e.g., increasing payments, redirecting payments to accounts the adversary controls), or creating or manipulating data to change payments to partners. Adversaries may also seek to acquire salable or fraudulently useful personally identifiable information (PII) about legitimate claimants (e.g., account numbers). They do this by causing cyber effects: modification or insertion of data used by the CRM application, modification of the application itself, or exfiltration or interception of claim or case data.

Adversaries can be insiders, partners, or external entities (e.g., criminals who have compromised end-user devices, or who have launched attacks against enterprise services). While human error is not a direct threat source, adversaries can take advantage of errors by end users, resources, administrators, or partners. Similarly, while natural disaster or another external event is not a direct threat source, adversaries can take advantage of an unexpectedly large claim load.

Adversaries are assumed to operate in a sustained timeframe (over months or even years), persistently planning and executing a cyber campaign. Adversaries' concern for stealth is assumed to be moderate, focused on concealing evidence of their presence, TTPs, and capabilities. For purposes of this use case, adversary targeting is assumed to be very narrow: adversaries target the CRM services, workflow, and data.

The threat model assumes that an adversary has established a presence on an enterprise system. In the threat scenarios considered in this use case, the adversary establishes their presence in the CRM application, by compromising a CRM endpoint in the presentation tier, leveraging the compromise of another application, or leveraging the compromise of a supporting enterprise service. The adversary expands their presence in the CRM application, by performing some combination of discovery, lateral movement, credential access, and privilege escalation. The adversary manages the system resources they have compromised, using command and control (C2) and evasion. The adversary uses TTPs as described in ATT&CK™ or the NSA/CSS Technical Cyber Threat Framework [12]. Because only a notional CRM application is identified in this use case, specific TTPs are not identified.

3.4 Alternative Solution

In this use case, only one alternative solution is considered: micro-segmentation. In the micro-segmentation solution, the organization will:

- Place claim filing, workflow automation, resource management, analytics, and reporting each in a separate segment (virtual enclave). This is referred to as "micro-segmentation" because segmentation is already applied to the CRM application in the three-tier architecture.
- Develop a new capability or configure existing IDS / insider threat tools to monitor interfaces or communications between segments.
- Develop a new capability or configure existing administrative / performance management tools to dynamically isolate or terminate suspect services. Note that performance gains may be obtained by managing segments separately, to respond to changes in demand.

3.5 Scores, Metrics, and MOEs

The CRM use case employs scoring, using SSM-CR, for its baseline assessment and the assessment of the identified solution. A descriptive metric is whether the solution could produce effects on adversary activities consistent with the organization's risk management strategy. Specific metrics which could be evaluated (in a modeling or emulation environment) can be identified from the Cyber Resiliency Metrics Catalog, based on the activities for which SME judgment determines that micro-segmentation would provide significant improvement.

3.5.1 Scoring: Baseline Assessment

To perform an assessment of the cyber resiliency capabilities of the existing CRM application, cyber resiliency constructs must be interpreted and prioritized for the situation described above. In this use case, the cyber resiliency constructs considered are objectives, sub-objectives, and activities. It must be emphasized that this prioritization is specific to the CRM application. Because the CRM application depends on enterprise services, many of the objectives are not relevant to – have zero priority for – the CRM application. An assessment for the enterprise infrastructure as a whole, or for specific enterprise services or functions would be quite different. For example, the Prevent / Avoid objective could be expected to be high or very high priority for the enterprise infrastructure; the weighting of the sub-objectives would reflect the assumption that the organization's SOC is relatively unsophisticated.

Table 1. Priority Weightings for Objectives and Sub-Objectives in the CRM Use Case

OBJECTIVE	OBJECTIVE PRIORITY WEIGHT	SUB OBJECTIVES	INTERPRETATION AND PRIORITY OF SUB OBJECTIVES FOR CRM APPLICATION
Prevent / Avoid Preclude the successful execution of an attack or the realization of adverse conditions.	0	• Apply basic cyber hygiene and risk-tailored controls. • Limit exposure to threat events. • Decrease the adversary's perceived benefits. • Modify configurations based on threat intelligence.	No change. The methods to achieve this objective are applied at the enterprise level.
Prepare Maintain a set of realistic courses of action that address predicted or anticipated adversity.	0	• Create and maintain cyber courses of action. • Maintain the resources needed to execute cyber courses of action. • Validate the realism of cyber courses of action. • Use validation methods that include testing or exercises.	No change. The methods to achieve this objective are applied at the enterprise level.
Continue Maximize the duration and viability of essential mission or business functions during adversity.	4	• Minimize degradation of service delivery. • Minimize interruptions in service delivery. • Ensure that ongoing functioning is correct.	Ensure that CRM services continue to be provided correctly and in a timely manner, despite adversity. • Minimize degradation of service delivery. (3) • Minimize interruptions in service delivery. (4) • Ensure that ongoing functioning is correct. (5)
Constrain Limit damage from adversity.	5	• Identify potential damage. • Isolate resources to limit future or further damage. • Move resources to limit future or further damage. • Change or remove resources and how they are used to limit future or further damage.	Limit damage from adversary modifications to or disruption of CRM services, behavior, or data. • Identify suspect claims and workflows. (5) • Isolate suspect claims and workflows, to ensure that disruption of processing for one claim (e.g., due to suspect data) does not cause cascading effects on other claims processing. (5) • Move services to limit future or further damage. (3)

OBJECTIVE	OBJECTIVE PRIORITY WEIGHT	SUB-OBJECTIVES	INTERPRETATION AND PRIORITY OF SUB-OBJECTIVES FOR CRM APPLICATION
Reconstitute Restore as much mission or business functionality as possible after adversity.	5	• Identify untrustworthy resources and damage. • Restore functionality. • Heighten protections during reconstitution. • Determine the trustworthiness of restored or reconstructed resources.	Restore correct and timely delivery of CRM services as quickly and completely as possible after discovery of and remediation of an incident. • Identify suspect services and data. (5) • Restore CRM service delivery. (5) (The last two methods are applied at the enterprise level.)
Understand Maintain useful representations of mission and business dependencies and the status of resources with respect to possible adversity.	0	• Understand adversaries. • Understand dependencies on and among systems containing cyber resources. • Understand the status of resources with respect to threat events. • Understand the effectiveness of cybersecurity and controls supporting cyber resiliency.	No change. The methods to achieve this objective are applied at the enterprise level.
Transform Modify mission or business functions and supporting processes to handle adversity and address environmental changes more effectively.	3	• Redefine mission / business process threads for agility. • Redefine mission / business functions to mitigate risks.	Modify CRM workflows to reduce risks due to adversarial activities. Ensure that oversight workflow processes are defined and executed effectively. (5) (The method related to agility does not apply.)
Re-Architect Modify architectures to handle adversity and address environmental changes more effectively.	5	• Restructure systems or subsystems to reduce risks. • Modify systems or subsystems to reduce risks.	Modify the CRM architecture to reduce risks. Restructure the CRM architecture for increased transparency, ease of detecting suspicious behavior, and greater responsive capability. (5)

An assessment of how well – how effectively, in light of the threat model – relevant activities are performed for the existing CRM applications produces an overall cyber resiliency score of 22.

3.5.2 Scoring: Assessment of Micro-Segmentation

Table 2 presents the summary of the assessments for the baseline and for micro-segmentation of the CRM application. These assessments use the scoring system illustrated in Figure 5. Micro-segmentation shows a significant improvement. It must be emphasized that these scores are situated in the assumptions identified above. Actual performance may be different, depending on how effectively the enterprise uses micro-segmentation and Analytic Monitoring (e.g., IDS, insider threat tools), as well as other cyber resiliency techniques (e.g., Non-Persistence, Substantiated Integrity). This assessment must not be confused with an assessment for the enterprise infrastructure, or for any specify enterprise service.

Table 2. Situated Cyber Resiliency Scoring for the CRM Use Case

Cyber Resiliency Performance Score			**Baseline: 22**	**Micro-Segmentation: 70**
Objective	*Restatement and Rationale for Priority Rating*	*Priority Rating for Objective*	*Achievement Score for Objective*	
Continue	Ensure that CRM services continue to be provided correctly and in a timely manner, despite adversity.	4	32	66
Constrain	Limit damage from adversary modifications to or disruption of CRM services, behavior, or data.	5	15	77
Reconstitute	Restore correct and timely delivery of CRM services as quickly and completely as possible after discovery of and remediation of an incident.	5	10	53
Transform	Modify CRM workflows to reduce risks due to adversarial activities.	3	60	80
Re-Architect	Modify the CRM architecture to reduce risks.	3	0	80

This solution is considered at a high level. More specific alternatives could be defined. For example, each virtual enclave may be able to incorporate Non-Persistence and Substantiated Integrity solutions. Segmentation could be extended further by decomposing each service. The potential cyber resiliency improvements of these more specific alternatives could be explored by identifying and evaluating cyber resiliency metrics.

3.5.3 Descriptive Metrics

One type of descriptive metric is the potential effects on adversary TTPs, as illustrated in Table 3. Because the system concept assumes a low-capability SOC, the risk management strategy focuses on the Preclude, Impede, and Limit effects. All Redirect effects, and the Scrutinize and Reveal aspects of Expose, are deprecated.

Table 3. Potential Effects of Micro-Segmentation of CRM Application on Adversary Activities

Effect	**Micro-Segmentation**
Redirect (includes deter, divert, and deceive)	Can support Divert and Deceive – Micro-segmentation provides opportunities for creating a deception environment.
Preclude (includes expunge, preempt, and prevent)	Can support Expunge (if Non-Persistence is applied to virtual enclaves)
Impede (includes contain, degrade and delay)	Contain (within a virtual enclave), Delay (takes longer to compromise multiple micro-services)
Limit (includes shorten and recover)	Shorten (due to detection), Recover (can recover an individual service, rather than needing to recover the entire CRM application)
Expose (includes detect, scrutinize and reveal)	Detect, can support Scrutinize and Reveal (depending on SOC capabilities and enterprise policies / practices regarding cyber threat intelligence information sharing)

Another descriptive metric consists of determining whether new attack scenarios are enabled by, or existing scenarios are made more likely by, the potential solution. Micro-segmentation increases the attack surface; if the use case is elaborated to include details about the CRM application, the resulting more detailed threat scenarios (including specific representative adversary TTPs) can be used to determine the extent to which those scenarios are made more likely by the increased attack surface. In addition, micro-segmentation can increase the potential for administrator error.

3.5.4 Quantitative Metrics

The following are examples of metrics which could be specified in detail and used to support (or disconfirm) the assessments of how well cyber resiliency-supporting activities could be performed without and with micro-segmentation:

- Percentage or number of detected adversary TTPs
- Time between compromise of CRM endpoint and detection of malicious activities
- Time between compromise of CRM endpoint and disruption of CRM functioning
- Time between compromise of CRM endpoint and fraudulent disbursement

As noted above, these metrics could be evaluated in a modeling or emulation environment, or in a testbed.

4 Vehicle Use Case

The vehicle use case, as summarized in Figure 8, involves the procurement of a fleet of vehicles, for use by enterprise staff within, near, and between enterprise campuses. These vehicles include a low degree of autonomy (e.g., driver assistance [16] [17]). Goals of cyber adversaries targeting a vehicle include manipulation of controls to cause an accident, physical theft, tracking enterprise staff movement, and eavesdropping on enterprise staff in the vehicle. The vehicle's cyber attack surface includes its embedded control units (ECUs), its entertainment system (which also manages operator displays), the Controller Access Network (CAN) bus, and its keyless entry system (KES). The acquisition program does not have the option of purpose-built vehicles, but can make changes or additions to commodity vehicles, including the integration of purpose-built components. However, such changes or additions must not preclude maintenance upgrades of the commodity vehicles (e.g., ECU software upgrades). For purposes of this use case, the focus is on a specific threat scenario, in which an adversary injects malicious commands via the infotainment system, to manipulate ECUs and cause an accident.

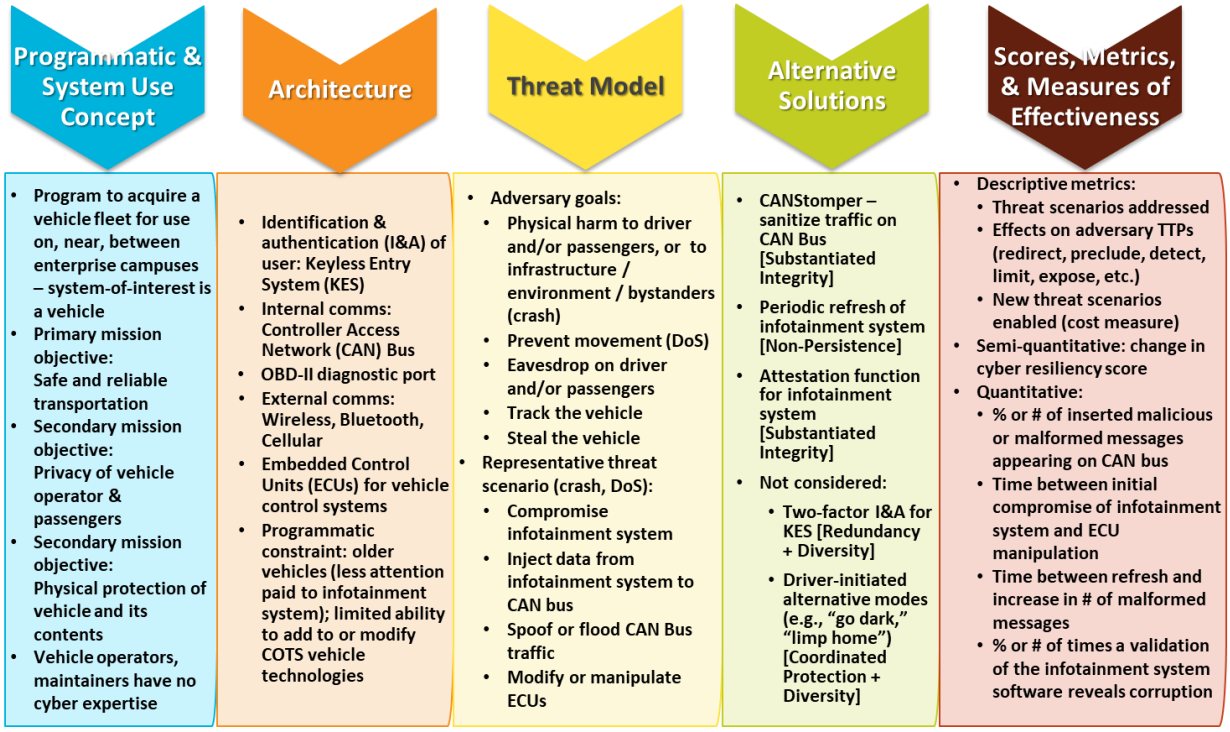

Figure 8. Summary of the Vehicle Use Case

4.1 Programmatic and System Use Concept

The system-of-interest is a vehicle in a fleet being acquired by an organization with multiple campuses in a region where the organization's staff might be targeted.

4.1.1 Background: Vehicle Cybersecurity and Cyber Resiliency Challenges

Automobiles are far more than a combustion engine and gearing. They include cyber-physical components such as sensors, ECUs, and fuel injectors, as well as purely cyber components such as an infotainment system, and connectivity to Bluetooth devices, cellular telephony, wireless networks, and satellite navigation and communication. The intertwined computer systems and embedded control units bring a large amount of functionality and flexibility to these previously mechanical only systems. The

modern vehicle is a real-time system with a deeply integrated computation core; it is a cyber-physical system.

As a cyber-physical system, the attack surface grows just as quickly as the flexibility and the capabilities do. Every new wireless connection yields another attack vector. Although computerization of the automobile means an increase in the safety from physical events, that doesn't transfer to safety from cyber threat events. Automatic braking and collision detection are feats of modern engineering. But modern engineering can fail to observe that "safety" means more than protection from physical events. Modern engineering needs to use and understand cyber resilience and security. Otherwise, the automatic breaking may not be braking for an imminent collision; automatic braking could be actuating for a keyboard miles away.

As with many cyber physical systems, the addition of cyber elements including wireless connectivity has been done to enhance the user experience, reduce costs, or in some instances enhance safety of the vehicle. It is only in recent years that vehicle manufacturers are recognizing that the addition of cyber elements can in some instances expose the vehicle to attack and that there are individuals and entities that have an interest in taking advantage of such exposures. While some manufacturers are beginning to attempt to take corrective action, the vast majority of vehicles are not built with security or resiliency in mind. Security of most vehicle's cyber components are largely via security through obscurity. In other words, the vehicle's CAN bus implicitly trusts any message sent on it and car manufacturers do not release the translation of arbitration identifiers (IDs, i.e., names of ECUs or functions) and the data payloads. But if an adversary can gain access to the CAN bus and steal or reverse engineer the arbitration IDs, they can gain control of vehicle components. Access to the CAN bus can be achieved in multiple ways, including compromise of a low-criticality sub-system such as the infotainment system. This state of the vehicle security and resiliency is the underlying assumption for the vehicle use case that follows.

4.1.2 Mission

For the purposes in this use case, the primary mission of the vehicle fleet is to provide safe and timely transportation of enterprise staff within and between campuses in a contested or problematic environment. Safety is very high criticality, while timeliness is high criticality. A secondary mission is to maintain the privacy of organization staff (e.g., geolocation, cell phone communications, conversations within the vehicle). The criticality of privacy protection is also high; organization staff may be targeted, and communications or in-vehicle conversations may reveal high-confidentiality organization information.

4.1.3 Environmental Assumptions

Two key technical assumptions are that

- The vehicle has a low level of automation, in terms of the levels defined by the Society of Automotive Engineers (SAE). The vehicle will have automation level 0 (e.g., no driver assistance except cruise control) or level 1 (e.g., adaptive cruise control, lane assist, in some cases parking assistance). [18] [19]
- A commodity infotainment system provides cellular communications, GPS (Global Positioning System) navigation, Bluetooth, Wi-Fi, and personalization of entertainment and user interface.

Even with a low level of automation, control messages on the CAN bus can instruct ECUs to increase or cut acceleration. Vulnerabilities in commodity infotainment systems are increasingly documented [20].

Assumptions about the operational environment include

- A moderate-sized population of organizational users (a.k.a. drivers), with no awareness of cyber threats to vehicles.

- A small population of vehicle maintenance staff, with little-to-no awareness of cyber threats to vehicles, and no responsibility for vehicle software maintenance beyond installing upgrades and patches when instructed by the organization's vehicle program.

In light of the operational assumptions, the decision environment is simple: a driver must decide whether or not to drive a vehicle, based on the information it presents via displays, and maintenance staff must decide whether to let a vehicle leave the vehicle fleet maintenance facility, based on the information they can obtain via diagnostics.

4.1.4 Programmatic Constraints

As noted above, the organization's vehicle fleet program seeks to acquire a commodity vehicle with minor enterprise modifications. A candidate vehicle has been identified, with low autonomy. That vehicle includes a commodity infotainment system. The organization will not seek to acquire a full telematics system, since privacy concerns mean that the organization does not want to transmit vehicle information to an insurer (or to a roadside assistance service, even if one is available in the region where the fleet will operate). The program is responsible for validating upgrades and patches from original equipment manufacturers (OEMs), making upgrades and patches to any vehicle modifications, and passing these on to maintenance staff.

4.2 Architecture

The basic architecture of the vehicle under consideration is shown in Figure 9. ECUs (e.g., engine, brakes, acceleration, temperature, fans, airbags) communicate with driver controls via the CAN bus.[6] The status of the vehicle (e.g., speed, temperature) is presented to the driver via displays, connected to the CAN bus and managed by the infotainment system. A low-assurance firewall (or gateway) is implemented between the infotainment system and the low-speed CAN bus; however, it is known to be vulnerable. Some manufacturers and third parties are implementing IDSs, but an IDS capability is not present in the vehicle under consideration.

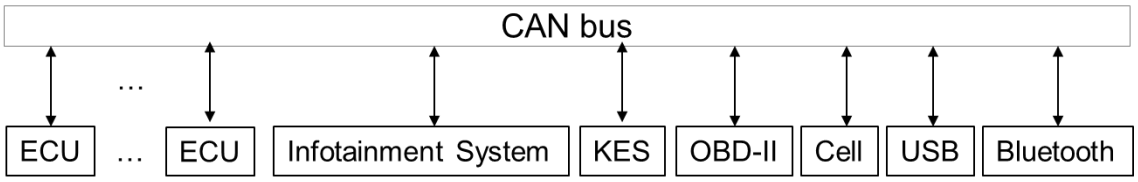

Figure 9. Vehicle Architecture

This architecture presents a varied attack surface to adversaries. Table 4 provides a representative set of attack vectors that an adversary could use to gain access to the CAN bus on the vehicle under consideration. (For a more extensive discussion, not limited to the vehicle under consideration, see [21] [22].) Those marked with an asterisk (*) require the adversary to bypass the firewall between the infotainment system and the CAN bus.

[6] In fact, the architecture has multiple CAN buses: ECUs for the engine, brakes, and acceleration are on one high-speed CAN bus, bridged to another high-speed CAN bus and to a low-speed CAN bus. However, because no firewalls exist between the buses, messages on one bus can move to another without restriction. Thus, the phrase "the CAN bus" is used for simplicity.

Table 4. Examples of Attack Vectors Against the CAN Bus

Attack Vector	Attack Range	Example of Published Discussion	Requirements on Attacker	Potential Impact
Cellular data connection	Long range, provides access to the internet and internet based services	[23]	Requires exploitation of the infotainment system or wireless protocol*	Can allow control of infotainment system, message injection onto the CAN bus
Tire Pressure Monitoring System (TPMS)	Short range wireless signal	[24]	Believed to be a very difficult attack vector, likely requires a buffer overflow	Can allow code execution, message injection onto the CAN bus
Bluetooth	Short range wireless signal	[25]	Requires exploitation of the infotainment system or wireless protocol*	Can allow control of infotainment system, message injection onto the CAN bus
Wi-Fi	Medium range	[23]	Requires exploitation of the infotainment system or wireless protocol*	Can allow control of infotainment system, message injection onto the CAN bus
Universal Serial Bus (USB)	Short range	[23]	Requires physical access and exploitation of the infotainment system*	Can allow control of infotainment system, message injection onto the CAN bus
Keyless Entry System (KES)	Short range	[26]	Requires exploitation of KES cryptographic system	Can allow unauthorized access to vehicle
OBD-II (On-Board Diagnostic System)	Short range	[27]	Requires physical access, possibly requires gateway bypass	Can allow full access to CAN bus, ability to modify ECUs
Exterior Sensors	Short to medium range	[28]	Requires physical exploitation or exploitation of the sensed medium	Can leverage factory "Limp Mode"[7] when sensor is altered or perturb sensor information

4.3 Threat Model

This use case focuses on an adversary with moderate-to-high knowledge of vehicle architectures and vulnerabilities. The adversary is assumed to have either physical access to the vehicle (e.g., while it is parked outside of the organization's campuses) or access to communications with vehicle systems (e.g., via cellular communications). (The threat model for a vehicle in general could be broadened to include access to vehicle systems outside of normal operations, e.g., during manufacture or maintenance. However, this is outside the scope of this use case.) Adversary goals include

[7] Limp mode is a protective function for a vehicle's engine and transmission, invoked when a faulty value from the engine or transmission control unit is detected. While functional details differ among vehicle manufacturers, limp mode typically limits engine revolutions per minute (RPM) and may restrict other functions. [30]

- Causing physical harm, by harming the driver or passengers (e.g., crashing the vehicle) or by harming the environment in which the vehicle operates (e.g., crashing the vehicle into pedestrians or property).
- Acquiring information that would enable the adversary to harm the organization, by tracking vehicles or drivers, eavesdropping on cell phone conversations, or eavesdropping on conversations within the vehicle.
- Interfering with organizational operations, by rendering vehicles undrivable.
- Personal gain, by stealing the vehicle.

The intended cyber effects include modification or insertion of instructions to or readings from ECUs, and interception of communications to or conversations within the vehicle. The adversary is assumed to operate over a sustained timeframe (e.g., months), persistently planning and executing a cyber campaign. The adversary's concern for stealth is limited, taking advantage of lack of organizational capabilities to detect malicious cyber activities against the vehicle fleet. The adversary's targeting is very narrow, focused on vehicles and communications rather than other organizational resources.

Threat scenarios involving physical harm or loss of confidentiality start in the same way: The adversary establishes and expands a presence on vehicle systems. To establish a presence, the adversary can gain physical access to vehicle and insert a device into the ODB-II port, compromise the infotainment system via a supply chain attack or a compromised end-user device (e.g., smartphone), compromise a type of ECU via a supply chain attack, or compromise the infotainment system via the attack vectors identified above. The adversary can take advantage of their presence on the CAN bus to compromise ECUs.

In representative threat scenarios which cause physical harm, the adversary manages compromised resources by performing command and control (C2) via wireless or cellular communications, or by establishing triggering conditions under which an effect will be caused (e.g., vehicle speed, location). The adversary can cause an effect by modifying or fabricating ECU sensor data to mislead the driver; modifying or fabricating commands to ECUs to direct vehicle behavior, causing the vehicle to brake, steer, or accelerate harmfully; or modifying or fabricating CAN bus traffic to cause denial-of-service (DoS) for selected vehicle functions (see ICS-ALERT-17-209-01).

In representative threat scenarios to gain information, the adversary can track the vehicle, performing C2 via wireless or cellular communications and causing covert cell or wireless communications, which reveal the vehicle's location. The adversary can use C2 via wireless or cellular communications to eavesdrop on cell phone communications, either sending a copy of cellular communications via wireless communications or recording cellular communications for later covert cellular transmission. The adversary can eavesdrop on conversations within vehicle by toggling the in-vehicle microphone via wireless or cellular communications, and sending covert cellular transmission or wireless communications to an adversary listening post.

4.4 Alternative Solutions

To identify alternative solutions, cyber resiliency constructs are interpreted and prioritized. Representative examples of risk mitigations which could be made part of a solution are identified. The focus of the use case presented here is then narrowed, for expository simplicity, to focus on threat scenarios involving physical harm, assuming a compromised infotainment system.

4.4.1 Interpretation and Prioritization of Cyber Resiliency Constructs

The cyber resiliency objectives and sub-objectives were interpreted and prioritized for the vehicle use case as shown in Table 5.

Table 5. Interpretation and Prioritization of Cyber Resiliency Objectives for Vehicle Use Case

OBJECTIVE	OBJECTIVE PRIORITY WEIGHT	SUB-OBJECTIVES	INTERPRETATION AND PRIORITY OF SUB-OBJECTIVES FOR CRM APPLICATION
Prevent / Avoid Preclude the successful execution of an attack or the realization of adverse conditions.	4	• Apply basic cyber hygiene and risk-tailored controls. • Limit exposure to threat events. • Decrease the adversary's perceived benefits. • Modify configurations based on threat intelligence.	Prevent access to, and resist interference with correct functioning of, vehicle systems. *High priority in operations. Accept the possibility that vehicle will not start if it has been improperly modified.* • Apply basic security engineering to the vehicle electronics; tailor controls to protect the most critical subsystems. (5) • Limit the exposure of ECUs to tampering and malicious instructions, particularly from the infotainment system. (5) (Last two methods deemed unrealistic for population of vehicle operators and maintenance staff.)
Prepare Maintain a set of realistic courses of action that address predicted or anticipated adversity.	0	• Create and maintain cyber courses of action. • Maintain the resources needed to execute cyber courses of action. • Validate the realism of cyber courses of action. • Use validation methods that include testing or exercises.	Provide operating procedures and supporting resources so that vehicle operator can respond to interference. *Not applicable. Deemed unrealistic for population of vehicle operators.*
Continue Maximize the duration and viability of essential mission or business functions during adversity.	5	• Minimize degradation of service delivery. • Minimize interruptions in service delivery. • Ensure that ongoing functioning is correct.	Ensure that vehicle operator can safely reach destination. *Top priority in operations.* • Enable the vehicle to operate correctly, even if some non-essential functions (e.g., radio) are disabled. (5) • Minimize the set of circumstances in which the vehicle does not operate at all. (5) • Ensure that critical vehicle systems are functioning within specifications and characteristic behavior. (4)

OBJECTIVE	OBJECTIVE PRIORITY WEIGHT	SUB-OBJECTIVES	INTERPRETATION AND PRIORITY OF SUB-OBJECTIVES FOR CRM APPLICATION
Constrain Limit damage from adversity.	4	• Identify potential damage. • Isolate resources to limit future or further damage. • Move resources to limit future or further damage. • Change or remove resources and how they are used to limit future or further damage.	Limit damage to primary systems; accept damage to or unavailability of non-essential systems and functioning. *High priority in operations. Accept the possibility that some systems (e.g., navigation, entertainment, communications) will become unavailable.* • Identify potential damage to, and circumstances which could result in harm to, critical systems.(5) • Isolate critical systems from malfunctioning or harmful non-essential systems. (5) (Last two methods deemed unrealistic, given architectural and programmatic constraints.)
Reconstitute Restore as much mission or business functionality as possible after adversity.	0	• Identify untrustworthy resources and damage. • Restore functionality. • Heighten protections during reconstitution. • Determine the trustworthiness of restored or reconstructed resources.	Enable maintenance staff to restore vehicle to an acceptable state. *High priority for maintenance – does not apply during operations. Currently deemed not applicable.*
Understand Maintain useful representations of mission and business dependencies and the status of resources with respect to possible adversity.	0	• Understand adversaries. • Understand dependencies on and among systems containing cyber resources. • Understand the status of resources with respect to threat events. • Understand the effectiveness of cybersecurity and controls supporting cyber resiliency.	Enable vehicle operator to understand the posture of vehicle systems with respect to cyber threats. Enable maintenance staff to determine whether, which, and how vehicle systems have been compromised. *Not applicable. Deemed unrealistic for vehicle operators, maintenance staff.*
Transform Modify mission or business functions and supporting processes to handle adversity and address environmental changes more effectively.	0	• Redefine mission / business process threads for agility. • Redefine mission / business functions to mitigate risks.	Define procedures which enable operators to handle unexpected vehicle behavior effectively, and provide training. *Not applicable. Deemed unrealistic for population of vehicle operators.*

OBJECTIVE	OBJECTIVE PRIORITY WEIGHT	SUB-OBJECTIVES	INTERPRETATION AND PRIORITY OF SUB-OBJECTIVES FOR CRM APPLICATION
Re-Architect Modify architectures to handle adversity and address environmental changes more effectively.	0	• Restructure systems or subsystems to reduce risks. • Modify systems or subsystems to reduce risks.	Modify design and implementation of vehicle systems (including ECUs, CAN bus, OBD-II) to reduce risks. Not applicable. Incompatible with procurement of commodity vehicles.

Table 6 provides representative examples of how cyber resiliency techniques could be interpreted in the context of the vehicle architecture.

Table 6. Representative Examples of How Cyber Resiliency Techniques Could Be Applied

Technique	Description	Application to Vehicle Systems
Adaptive Response	Implement nimble cyber courses of action (CCoAs) to manage risks	Make changes to ECUs while they continue operating. Ideally this would be in response to out-of-bounds behavior.
Coordinated Protection	Ensure that protection mechanisms operate in a coordinated and effective manner	Coordinate and correlate data from bus traffic and external input to detect out of baseline behavior. Utilize multiple factors of authentication before input specific actions are taken.
Privilege Restriction	Restrict privileges based on attributes of users and system elements as well as on environmental factors	Define, assign, maintain, and apply strict restrictions on users accessing the vehicle via wireless protocols.
Segmentation	Define and separate system elements based on criticality and trustworthiness	Define a critical enclave that is made up of the key ECUs for basic vehicle function. Separate this enclave from normal communication buses with the ability to take control given physical user affirmation.
Substantiated Integrity	Ascertain whether critical system elements have been corrupted	Apply and validate checks of the integrity or quality of the ECUs and information on the CAN bus.

Based on programmatic constraints and the overall risk management strategy, the strategic cyber resiliency design principles were assessed for relevance; representative examples of their application are given in Table 7.

Table 7. Relevance of Strategic Cyber Resiliency Design Principles to the Vehicle Use Case

Strategic Cyber Resiliency Design Principle	Relevance	Application to Vehicle Use Case
Focus on Common Critical Assets	5	The CAN bus, and CAN bus traffic to and from ECUs, are critical common assets. Examples of application include: Coordinate and correlate data from bus traffic and external input to detect out of baseline behavior. Utilize multiple factors of authentication before input specific actions are taken.
Support and Architect for Extensibility	0	Not applicable (N/A)
Reduce the Attack Surface	2	The attack surface includes all attack vectors against the CAN bus and its traffic. Representative examples are given in Table 4. (The attack surface also includes the supply chain and the maintenance environment; because these are out of scope, the relevance of this design principle is relatively low.) Examples of application include: Strengthen or eliminate APIs that are vulnerable to cyber-attack. Apply the principle of least privilege to message traffic.
Assumed Compromised Resources	3	Resources that could be compromised include ECUs as well as the infotainment system; because this use case focuses on a compromised infotainment system only, the relevance of this design principle is only moderate. Examples of application include: Validate the correctness of the critical ECUs by comparing configurations with a baseline digital signature. Validate the correctness of CAN bus messages by comparing messages to baseline behavior.
Expect Adversaries to Evolve	0	N/A

Based on the architecture, programmatic constraints, and the relevance of strategic design principles, the structural cyber resiliency design principles were assessed for relevance to this use case. Representative examples of how each design principle could be applied (in general, not simply in this use case) are provided in Table 8.

Table 8. Relevance of Structural Design Principles to the Vehicle Use Case

Structural Principles	Relevance	Example of Application to Vehicle Systems in General
Limit the need for trust	0	Limit privileges to critical ECUs. Create an enclave of critical ECUs as described in the Segmentation technique above.
Control visibility and use	2	Disable non-critical elements displaying suspicious behavior. Limit privileges to critical ECUs.
Contain and exclude behaviors	5	Validate the correctness of the critical ECUs by comparing configurations with a baseline digital signature. Validate the correctness of CAN bus messages by comparing messages to baseline behavior.
Layer and partition defenses	3	Create an enclave of critical ECUs as described in the Segmentation technique above.
Plan and manage diversity	0	The car manufacturer acquires critical ECUs from multiple suppliers.
Maintain redundancy	0	Create a secondary set of ECUs that can take over if systems are not functioning properly.
Make resources location versatile	0	N/A

Structural Principles	Relevance	Example of Application to Vehicle Systems in General
Leverage health and status data	5	Use data and correlate data to behavior to determine state of vehicle.
Maintain situation awareness	0	N/A
Manage resources (risk) adaptively	3	Be able to shut down or negate compromised systems. Enable limp mode.
Maximize transience; minimize persistence	2	Periodically refresh ECU configurations to a baseline setting. Periodic refresh of the infotainment system.
Determine ongoing trustworthiness	2	Validate the integrity of data being communicated and the behavior of services.
Change or disrupt the attack surface	0	Shut off communication unless actively being used by vehicle.
Make unpredictability and deception user-transparent	0	N/A

4.4.2 Examples of Potential Mitigations

A wide variety of risk mitigations can be identified to improve the cyber resiliency of individual vehicles, for the organization's vehicle fleet. Examples are provided in Table 9, and are mapped to the cyber resiliency techniques and structural design principles they apply. Most of these are not relevant to this use case, given programmatic constraints, but are included for illustrative purposes.

Table 9. Examples of Potential Mitigations to Reduce Cyber Risk to Vehicles

Potential Mitigation	Description	Techniques Applied	Structural Design Principles Applied
CANStomper	Perform message source validation and suppression of invalid-source messages	Substantiated Integrity	Contain and exclude behaviors Leverage health and status data Determine ongoing trustworthiness
Re-flash infotainment system	Periodic refresh of the infotainment system	Non-Persistence	Maximize transience; minimize persistence
Validate infotainment system	Attestation function for the infotainment system	Substantiated Integrity	Determine ongoing trustworthiness
Firewall wireless devices	Create rules that dictate the type of traffic allowed via the wireless interfaces	Coordinated Protection Segmentation	Layer and partition defenses
Validate ECUs	Attestation function for ECUs	Substantiated Integrity	Determine ongoing trustworthiness
Intrusion Detection System (IDS)/Intrusion Protection System (IPS) for wireless communication	Utilize pattern recognition software to stop/disable malicious behavior from wireless interfaces	Analytic Monitoring	Assume compromised resources Determine ongoing trustworthiness
IDS/IPS for control layer (CAN) communication	Utilize pattern recognition software to stop/disable malicious behavior in the CAN bus	Coordinated Protection	Layer and partition defenses

Potential Mitigation	Description	Techniques Applied	Structural Design Principles Applied
Safe Mode non-essential ECU functionality	Force car into limited functioning mode which eliminates all unnecessary communication and operation.	Adaptive Response Coordinated Protection	Manage resources adaptively Focus on common critical assets Layer and partition defenses
"Go Dark" mode	Disable all wireless interfaces	Adaptive Response Coordinated Protection	Manage resources adaptively Focus on common critical assets Layer and partition defenses
Hood Lock	A secondary lock that prevents the hood from being opened without a key	Coordinated Protection	Layer and partition defenses
Digital Signing for ECU updates	ECU firmware can only be updated if it has the OEM digital signature	Substantiated Integrity	Determine ongoing trustworthiness
OBD-II lock	A physical or logical wall that prevents access to the CAN bus via the OBD-II unless authenticated	Coordinated Protection	Layer and partition defenses
Multi-Factor Authentication (MFA) for door lock	Door locks can only be digitally opened once an additional means of authentication has been passed	Diversity Coordinated Protection	Manage resources adaptively Plan and manage diversity

4.4.3 Alternative Solutions

The following solutions to limit potential damage from compromised infotainment system were identified:

- Insert a new component on CAN bus to perform message source validation and suppression of invalid-source messages. For purposes of this use case, the CANStomper prototype was considered; other versions have been researched elsewhere.
- Add a purpose-built periodic refresh function to the infotainment system.
- Add a purpose-built attestation function to the infotainment system.

Additional solutions related to vulnerabilities in Keyless Entry System (KES) or repurposing the different CAN buses were deemed out of scope.

The first alternative applies the Segmentation cyber resiliency technique and the structural design principles "Control visibility and use" and "Contain and exclude behavior." The second and third alternatives apply Substantiated Integrity (the periodic refresh alternative also applies Non-Persistence) and the "Determine ongoing trustworthiness" design principle.

4.5 Scores, Metrics, and MOEs

The vehicle use case employs scoring, using SSM-CR, for its baseline assessment and for assessments of alternative solutions. A descriptive metric is whether the solution could produce effects on adversary activities consistent with the organization's risk management strategy. Specific metrics which could be

evaluated (in a modeling or emulation environment) can be identified from the Cyber Resiliency Metrics Catalog, based on the activities for which SME judgment determines that the identified solutions would provide significant improvement.

4.5.1 Scoring

SME assessments of activities supporting achievement of relevant sub-objectives and objectives for the baseline and for the alternative solutions are presented in Appendix B. The summary of the results of the assessments for objectives are presented in Table 10. These results must be understood to be situated in a narrow threat model – considering only threat scenarios in which compromise of the infotainment system is exploited to cause physical harm to the vehicle, its driver and passengers, and the physical environment – and in the context of the organization's operational concept for and programmatic constraints on the vehicle fleet.

Table 10. Situated Cyber Resiliency Scoring for Vehicle Use Case

Cyber Resiliency Performance Score		Baseline 24	CANStomper 38	Refresh 29	Attestation 28	CANStomper + Refresh + Attestation 47
Objective	Priority Rating for Objective	Achievement Score for Objective				
Prevent / Avoid	4	24	34	42	24	48
Prepare	0	0	0	0	0	0
Continue	5	28	41	28	33	45
Constrain	4	17	39	17	24	46
Reconstitute	0	0	0	0	0	0
Understand	0	0	0	0	0	0
Transform	0	0	0	0	0	0
Re-Architect	0	0	0	0	0	0

4.5.2 Descriptive Metrics

One type of descriptive metric is the potential effects of alternative solutions on adversary TTPs, as illustrated in Table 11. Because the system concept assumes no end-user cyber expertise, and minimal awareness on the part of maintenance staff, the risk management strategy focuses on the Preclude and Impede effects. All Redirect effects, and the Scrutinize and Reveal aspects of Expose, are explicitly out of scope.

Table 11. Potential Effects of Alternatives in Vehicle Use Case on Adversary Activities

Effect	CANStomper	Refresh	Attestation
Redirect (includes deter, divert, and deceive)			
Preclude (includes expunge, preempt, and prevent)	Prevent (keep bad messages off the CAN bus)	Expunge	Expunge or Preempt (indirectly, as a consequence of detection)
Impede (includes contain, degrade and delay)	Contain, Degrade (reduce effectiveness of message flooding attacks)	Degrade	Contain (indirectly, as a consequence of detection)
Limit (includes shorten and recover)		Shorten	Shorten (indirectly, as a consequence of detection)
Expose (includes detect, scrutinize and reveal)			Detect

No new threat scenarios were identified as a result of the identified potential solutions.

4.5.3 Quantitative Metrics

The general process for using quantitative metrics to support analysis of alternatives is as follows: First, for each alternative solution, a *preliminary assessment* is made using SSM-CR, as shown above. Systems engineers determine which activities are expected to be performed more effectively, and thus which sub-objectives and objectives are expected to be better achieved.

In the example, CANStomper is expected to improve the performance of the high-priority "Restrict behaviors of users and cyber entities (e.g., components, services, processes, interfaces) based on degree of trust" activity, which supports achieving the top-priority "Apply basic hygiene and risk-tailored controls" sub-objective of the high-priority Prevent / Avoid objective, from very low to high. The rationale for this change is captured in the scoring worksheet: "CANStomper source arbitration stops the infotainment system from causing ECUs to act. However, it does not prevent the infotainment system from modifying a legitimate message." Similarly, potential performance of other activities under the Prevent / Avoid objective are improved, so that CANStomper changes the level of achievement for the Prevent / Avoid objective from 24 to 38.

Second, *potential cyber resiliency MOEs* are identified. Ideally, these should apply equally well across all solutions, to support comparison. MOEs can be identified in several ways:

- Potential metrics which serve as evidence for performance of an activity can be identified from the catalog. For example, one metric for the "Enforce clear boundaries on sets of cyber resources" activity is "Percentage of mission-critical cyber resources which can be discovered or reached from each enclave, sub-system, or network nodes." That metric can be tailored for the use case as "Percentage of ECUs which can be identifiably addressed from the entertainment system." Note that, depending on the evaluation environment and data-gathering tools available, it may be infeasible to evaluate the tailored metric directly.

- Mission MOEs which are expected to improve can be identified from the system concept of use, in the context of the threat scenario of concern. For example, if the vehicle operator steps on the brake, the vehicle is expected to slow. If the adversary has successfully executed the attack, inserting false commands to the braking system from the entertainment system, the vehicle will not slow; if the alternative solution successfully thwarts the attack, the vehicle should slow.
- Observable effects on adversary activities can be identified, based on specification of the threat scenario of concern. For example, in the scenario in which the adversary uses the entertainment system to send false commands to the braking system, one MOE is whether adversary-injected commands appear on the CAN bus.

Table 12 identifies a few representative examples of potential metrics for the different alternatives and the activities for which changes in metric values constitute evidence of effectiveness. (This list is by no means exhaustive.)

Table 12. Examples of Possible Metrics, Mapped to Activities, for Different Alternatives

Metric	Activity / Capability	Alternative(s)
# or % of messages that appear on CAN bus from infotainment system	PA-S1-A2: Restrict behaviors of users and cyber entities (e.g., components, services, processes, interfaces) based on degree of trust PA-S1-A5: Protect data in different states (e.g., at rest, in transit, in processing)	Refresh CANStomper
Time between refresh and when number of malformed (or total number) of messages increases again	PA-S1-A2: Restrict behaviors of users and cyber entities (e.g., components, services, processes, interfaces) based on degree of trust PA-S1-A5: Protect data in different states (e.g., at rest, in transit, in processing)	Refresh
% or number of times a validation of the infotainment system software reveals corruption	PA-S1-A5: Protect data in different states (e.g., at rest, in transit, in processing) CN-S3-A3: Validate software / service integrity / behavior to ensure it has not been corrupted	Refresh Attestation
Time between compromise of infotainment system and detection or notification	CS-S1-A2: Identify potentially compromised or faulty processes or services (i.e., those which can no longer be trusted)	Attestation
# or % of malformed messages that appear on CAN bus from the infotainment system	PA-S1-A2: Restrict behaviors of users and cyber entities (e.g., components, services, processes, interfaces) based on degree of trust	CANStomper
Percentage of malicious messages injected from the infotainment system appearing on the CAN bus	PA-S1-A5: Protect data in different states (e.g., at rest, in transit, in processing)	CANStomper

Third, the potential cyber resiliency MOEs are *downselected* based on the feasibility of evaluation. The selection takes into consideration whether the evaluation will be performed conceptually, in a model-based environment, in a laboratory or testbed, or in a representative operational environment. The selection also takes into consideration the time and effort required to evaluate each metric. In the vehicle use case, evaluation in a testbed is assumed to be feasible.

Fourth, the selected cyber resiliency MOEs are *evaluated*. The evaluation can be a single instance (as when a Red Team executes an attack), can involve repeated evaluation efforts (as when multiple runs of a simulation are generated), or can be exhaustive (involving all possible combinations of inputs, via simulation). The trade-offs between single-run, multiple-run, and exhaustive evaluations are informed by such factors as time, effort, engineering expectations of the solutions, and availability of measurement tools. For example, in a testbed environment for the vehicle use case, a tool for monitoring and analyzing traffic on the CAN bus is essential to evaluating the MOE of whether (or what percentage of) adversary-injected commands appear on the CAN bus. The observed behavior of the vehicle can be used not only to evaluate the mission MOE, but via experimentation can be used to estimate the percentage of addressable ECUs.

Fifth, the scores for relevant activities are *adjusted* based on the results of the evaluation(s) of selected MOEs, which provide evidence supporting the assignment of activity performance scores. These score adjustments ripple to changes in the scores for sub-objectives, objectives, and overall cyber resiliency. This enables comparison of the alternatives to determine which offer the greatest degree of cyber resiliency improvement, and whether the degree of cyber resiliency improvement offered by any alternative suffices.

5 References

[1] NIST, "Draft NIST Special Publication 800-160 Volume 2, Systems Security Engineering: Cyber Resiliency Considerations for the Engineering of Trustworthy Secure Systems," 21 March 2018. [Online]. Available: https://csrc.nist.gov/CSRC/media/Publications/sp/800-160/vol-2/draft/documents/sp800-160-vol2-draft.pdf.

[2] NIST, "NIST SP 800-160, Systems Security Engineering: Considerations for a Multidisciplinary Approach in the Engineering of Trustworthy Secure Systems," 15 November 2016. [Online]. Available: http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-160.pdf.

[3] D. J. Bodeau and R. D. Graubart, "Cyber Resiliency Design Principles: Selective Use Throughout the Lifecycle and in Conjunction with Related Disciplines," January 2017. [Online]. Available: https://www.mitre.org/sites/default/files/publications/PR%2017-0103%20Cyber%20Resiliency%20Design%20Principles%20MTR17001.pdf.

[4] D. Bodeau, R. Graubart, W. Heinbockel and E. Laderman, "Cyber Resiliency Engineering Aid - The Updated Cyber Resiliency Engineering Framework and Guidance on Applying Cyber Resiliency Techniques, MTR140499R1, PR 15-1334," May 2015. [Online]. Available: http://www.mitre.org/sites/default/files/publications/pr-15-1334-cyber-resiliency-engineering-aid-framework-update.pdf or http://www.defenseinnovationmarketplace.mil/resources/20150527_Cyber_Resiliency_Engineering_Aid-Cyber_Resiliency_Techniques.pdf.

[5] D. Bodeau and R. Graubart, "Cyber Resiliency Assessment: Enabling Architectural Improvement (MTR 120407, PR 12-3795)," May 2013. [Online]. Available: http://www.mitre.org/sites/default/files/pdf/12_3795.pdf.

[6] D. Bodeau and R. Graubart, "Cyber Resiliency Engineering Framework (MTR110237, PR 11-4436)," September 2011. [Online]. Available: http://www.mitre.org/sites/default/files/pdf/11_4436.pdf.

[7] D. Bodeau, R. Graubart, R. McQuaid and J. Woodill, "Cyber Resiliency Metrics, Measures of Effectiveness, and Scoring: Enabling Systems Engineers and Program Managers to Select the Most Useful Assessment Methods (MTR 180314)," The MITRE Corporation, Bedford, MA, 2018.

[8] D. Bodeau, R. Graubart, R. McQuaid and J. Woodill, "Cyber Resiliency Metrics Catalog (MTR 180450)," The MITRE Corporation, Bedford, MA, 2018.

[9] CPS PWG, "Framework for Cyber-Physical Systems, Release 1.0," May 2016. [Online]. Available: https://s3.amazonaws.com/nist-sgcps/cpspwg/files/pwgglobal/CPS_PWG_Framework_for_Cyber_Physical_Systems_Release_1_0Final.pdf.

[10] DoD, "DoD Program Manager's Guidebook for Integrating the Cybersecurity Risk Management Framework (RMF) into the System Acquisition Lifecycle, Version 1.0," 26 May 2015. [Online]. Available: https://acc.dau.mil/adl/en-US/722603/file/80119/Cybersecurity%20Guidebook%20v1_0%20with%20publication%20notice.pdf.

[11] NIST, "Framework for Improving Critical Infrastructure Cybersecurity, Version 1.1," 16 April 2018. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.04162018.pdf.

[12] National Security Agency, "NSA/CSS Technical Cyber Threat Framework v1," 6 March 2018. [Online]. Available: https://www.iad.gov//iad/library/reports/assets/public/upload/NSA-CSS-Technical-Cyber-Threat-Framework-v1.pdf.

[13] The MITRE Corporation, "Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK™)," The MITRE Corporation, 2015. [Online]. Available: https://attack.mitre.org/wiki/Main_Page.

[14] D. Bodeau and R. Graubart, "Cyber Prep 2.0: Motivating Organizational Cyber Strategies in Terms of Preparedness (MTR 150264, PR 16-0939)," The MITRE Corporation, Bedford, MA, 2016.

[15] D. J. Bodeau, C. D. McCollum and D. B. Fox, "Cyber Threat Modeling: Survey, Assessment, and Representative Framework (PR 18-1174)," The MITRE Corporation, McLean, VA, 2018.

[16] Autonomy Priority Steering Council, "Brief on Autonomy Initiatives in the US DoD," 8 November 2012. [Online]. Available: http://www.defenseinnovationmarketplace.mil/resources/Autonomy-PSC_Briefing_DistroA_RE.pdf.

[17] A. Lacher, "Autonomy & Transportation: Addressing Cyber-Resiliency Challenges, PR 15-1841," 10 June 2015. [Online]. Available: http://csrc.nist.gov/groups/SMA/ispab/documents/minutes/2015-06/ispab_june-10_alacher.pdf.

[18] National Highway Traffic Safety Administration, "Automated Vehicles for Safety," [Online]. Available: https://www.nhtsa.gov/technology-innovation/automated-vehicles-safety.

[19] Car and Driver, "Path to Autonomy: Self-Driving Car Levels 0 to 5 Explained," Car and Driver, October 2017. [Online]. Available: https://www.caranddriver.com/features/path-to-autonomy-self-driving-car-levels-0-to-5-explained-feature.

[20] S. Mazloom, M. Rezaeirad and A. Hunter, "A Security Analysis of an In Vehicle Infotainment and App Platform," in *10th USENIX Workshop on Offensive Technologies (WOOT '16)*, Austin, TX, 2016.

[21] A. Oyler and H. Saiedian, "Security in automotive telematics: a survey of threats and risk mitigation strategies to counter the existing and emerging attack vectors," *Security and Communications Networks,* vol. 9, p. 4330–4340, 2016.

[22] C. Valasek and C. Miller, "A Survey of Remote Automotive Attack Surfaces," 26 September 2014. [Online]. Available: https://ioactive.com/pdfs/IOActive_Remote_Attack_Surfaces.pdf.

[23] A. Greenberg, "The Jeep Hackers Are Back to Prove Car Hacking Can Get Much Worse," 01 August 2016. [Online]. Available: https://www.wired.com/2016/08/jeep-hackers-return-high-speed-steering-acceleration-hacks/.

[24] I. Rouf, R. Miller, H. Mustafa, T. Taylor, S. Oh, W. Xu, M. Gruteser, W. Trappe and I. Seskar, "Security and Privacy Vulnerabilities of In-Car Wireless Networks: A Tire Pressure Monitoring System Case Study," in *Proceedings of the 19th USENIX conference on Security (USENIX'10)*, Washington, DC, 2010.

[25] A. Kovelman, "A Remote Attack on the Bosch Drivelog Connector Dongle," Argus Security, April 2017. [Online]. Available: https://argus-sec.com/remote-attack-bosch-drivelog-connector-dongle/.

[26] R. Benadjila, M. Renard, J. Lopes-Esteves and C. Kasmi, "One Car, Two Frames: Attacks on Hitag-2 Remote Keyless Entry Systems Revisited," in *11th USENIX Workshop on Offensive Technologies (WOOT '17)*, Vancouver, BC, Canada, 2017.

[27] D. Klinedinst and C. King, "On Board Diagnostics: Risks and Vulnerabilities of the Connected Vehicle," March 2016. [Online]. Available: https://resources.sei.cmu.edu/asset_files/WhitePaper/2016_019_001_453877.pdf.

[28] A. Greenberg, "Hackers Fool Tesla S's Autopilot to Hide and Spoof Obstacles," 4 August 2016. [Online]. Available: https://www.wired.com/2016/08/hackers-fool-tesla-ss-autopilot-hide-spoof-obstacles/.

[29] CNSS, "Committe on National Security Systems (CNSS) Glossary (CNSS Instruction No. 4009)," 26 April 2015. [Online]. Available: https://www.cnss.gov/CNSS/openDoc.cfm?hldYMe6UHW4ISXb8GFGURw==.

[30] J. S., "Limp Mode – Meaning, Causes & Solutions," Mechanic Base, 2018 15 January. [Online]. Available: https://mechanicbase.com/engine/limp-mode/.

Appendix A Details of Vehicle Use Case Scoring □

Table 13 presents details of the scoring for the baseline and alternatives in the vehicle use case. For brevity, only those objectives, sub-objectives, and activities which were determined to be relevant (non-zero priority rating) are included, and only changes in activity performance scores are indicated. Thus, for example, only the baseline performance score is given for the first activity under the "Apply basic hygiene and risk-tailored controls" sub-objective of Prevent / Avoid; this reflects the fact that none of the alternatives considered changed the expected performance of that activity.

Table 13. Details of Vehicle Use Case Scoring

Prevent / Avoid				
Apply basic hygiene and risk-tailored controls	*Sub-Objective Priority Rating:* 5	*Restatement / Rationale for Priority Rating:* Apply basic security engineering to the vehicle electronics; tailor controls to protect the most critical subsystems.	*Sub-Objective Score:* Baseline: 16 CANStomper: 34 Refresh: 31 Attestation: 16	
Activity	*Activity Priority Rating*	*Restatement / Rationale for Priority Rating*	*Activity Performance Score*	*Rationale*
Restrict access to resources based on criticality and sensitivity (i.e., on resource attractiveness to adversaries)	5	Protecting critical resources - ECUs and the paths to the ECUs - can stop adversaries from causing the vehicle to enter an unacceptable state.	Baseline: 1	Some ECUs are "locked" from being modified, and some paths are restricted, but this has been circumvented in the wild and in the research community.
Restrict behaviors of users and cyber entities (e.g., components, services, processes, interfaces) based on degree of trust	4	Acceptable behaviors of vehicle subsystems can be defined (not based on conventional definition of trust).	Baseline: 1 CANStomper: 4 Refresh: 3	ECUs will not generally act upon unknown commands. CANStomper source arbitration stops the infotainment system from causing ECUs to act. However, it does not prevent the infotainment system from modifying a legitimate message. Although infotainment system is not critical or sensitive, its periodic refresh eliminates / reduces effectiveness of adversary access to CAN bus and hence to ECUs. Performance depends on frequency of refresh.
Enforce clear boundaries on sets of cyber resources	5	Minimize contact between critical and noncritical ECUs	Baseline: 1	Firewall at diagnostic port. Firewall on entertainment system. Some manufacturers allow checking of state of the vehicle and not allowing diagnostic commands, but not the one under consideration. Some vehicles segment the CAN bus, but not the one under consideration.

Apply multiple defenses to critical assets	5	Protect critical ECUs with multiple layers of defenses (e.g., message traffic filtering).	Baseline: 0 CANStomper: 2	Not done CANStomper adds a layer of defense.
Protect data in different states (e.g., at rest, in transit, in processing)	5	Protect the integrity of the software on the ECUs. Protect the availability and quality of data on the CAN bus, or the vehicle will be inoperable.	Baseline: 1 Refresh: 3	Some basic check bits (a cyclic redundancy check or CRC) are used on messages on the CAN bus. Some handshake is needed to update ECUs, but can easily be compromised in a maintenance environment. Periodic refresh restores the integrity of infotainment system software and configuration data.
Limit exposure to threat events	*Sub-Objective Priority Rating:* 5	*Restatement / Rationale for Priority Rating:* Limit the exposure of ECUs to tampering and malicious instructions, particularly from the entertainment system.	*Sub-Objective Score:* Baseline: 33 CANStomper: 33 Refresh: 54 Attestation: 33	
Activity	*Activity Priority Rating*	*Restatement / Rationale for Priority Rating*	*Activity Performance Score*	*Rationale for Performance Score*
Define and implement a set of change parameters (e.g., conditions under which changes should not be made, "distance" beyond which a service should not be moved, ranges for frequency of changes)	5	Characterize system and component behaviors as "good" vs. "bad" (e.g., shutting off a critical system such as the fuel pump when the car is in motion is "bad" / unacceptable).	Baseline: 3	Many critical subsystems have defined conditions under which specific actions can or cannot be taken, based on safety concerns and the concern for protecting the engine against potential harm. "Limp home" mode is defined for the vehicle under consideration.
Switch to an alternative resource randomly or in response to a triggering event	3	Switch manually to secondary or alternative subsystem, based on operator observation and action. *Moderate priority, since operator action can be problematic.*	Baseline: 1	Operator can switch from primary braking system to physical handbrake (not an electronic handbrake). Other systems (CAN bus, steering, acceleration) are theoretical.
Retain resources in an active or "live" state for a limited lifespan (e.g., maximum time period after instantiation or creation, maximum period after use)	2	Can be applied to over-the-air updates, safety-critical communications (e.g., call 911). Some systems must stay on, even if the car is turned off (e.g., access, ignition).	Baseline: 0 Refresh: 4	For some vehicles, can enable over-the-air updates, but not for vehicle under consideration. For some vehicles, can power-on remotely (e.g., via OnStar), but not for vehicle under consideration. Infotainment system software is retained for a limited period.

Ensure that termination, deletion, or movement does not leave residual data or software behind	2	Relevant to fleet purchase. Applies to information stored by infotainment system, in a vehicle used serially by different users.	Baseline: 0 Refresh: 5	Not done in vehicle under consideration. Upon restart, infotainment system software is refreshed and configuration / user data is refreshed. Does not provide this capability for critical systems.
Separate cyber resources based on criticality and/or sensitivity	5	Want to separate (virtually as well as physically) safety-critical from non-critical systems.	Baseline: 2	For vehicle under consideration, engine, brakes, acceleration on one high-speed CAN bus, bridged to another high-speed CAN bus and to a low-speed CAN bus. No firewall between buses. Firewall (gateway) between infotainment system and low-speed CAN bus.
Decrease the adversary's perceived benefits	*Sub-Objective Priority Rating*: 0	*Restatement / Rationale for Priority Rating:* N/A. Deemed unrealistic for population of vehicle operators and maintenance staff.	0	
Modify configurations based on threat intelligence	*Sub-Objective Priority Rating:* 0	*Restatement / Rationale for Priority Rating:* N/A. Deemed unrealistic for population of vehicle operators and maintenance staff.	0	
Continue				
Minimize degradation of service delivery	*Sub-Objective Priority Rating:* 5	*Restatement / Rationale for Priority Rating:* Enable the vehicle to operate correctly, even if some non-essential functions (e.g., radio) are disabled.	*Sub-Objective Score:* Baseline: 40 CANStomper: 47 Refresh: 40 Attestation: 40	
Activity	*Activity Priority Rating*	*Restatement / Rationale for Priority Rating*	*Activity Performance Score*	*Rationale for Performance Score*
Perform mission damage assessment	5	Vehicle can self-evaluate (e.g., evaluate whether engine timing is within acceptable range) to determine whether a limited-functionality ("limp home") mode is warranted, or whether the operator must be notified (e.g., "check engine" light).	Baseline: 2	Not currently implemented comprehensively (only look at a few parameters, evaluate a few simple tests) or with sufficient consideration of operational environment (e.g., vehicle goes into "limp home" mode when traction is inadequate).
Maintain acceptable levels of performance for mission-critical, security-critical, and mission-supporting applications and services	5	Maintain acceptable levels of mission-critical functions - i.e., engine, steering, brakes, acceleration - so that the mission can be completed - i.e., the vehicle arrives safely. (Rely on mission damage assessment.)	Baseline: 2	Vehicle can go into "limp home" mode, with minimal damage to systems.

Select and tailor CCoA	5	CCoAs in the vehicle are purely automated; the vehicle operator is unaware of and indifferent to the cyber posture of vehicle systems. CCoAs in the vehicle involve turning off non-critical systems.	Baseline: 2 CANStomper: 3	The range of CCoAs in vehicles is currently limited to a few systems beyond "limp home" mode. For example, if traction control is assessed as damaged, it can be automatically disabled; if power steering fails, a fail-safe mode for steering engages. CANStomper removes messages from unidentified or illegitimate sources from the CAN bus.
Minimize interruptions in service delivery	*Sub-Objective Priority Rating:* 5	*Restatement / Rationale for Priority Rating:* Minimize the set of circumstances in which the vehicle does not operate at all.	*Sub-Objective Score:* Baseline: 31 CANStomper: 38 Refresh: 31 Attestation: 31	
Activity	*Activity Priority Rating*	*Restatement / Rationale for Priority Rating*	*Activity Performance Score*	*Rationale for Performance Score*
Perform mission damage assessment	5	Vehicle can self-evaluate (e.g., evaluate whether engine timing is within acceptable range) to determine whether the vehicle is safe to operate.	Baseline: 2	Not currently implemented comprehensively (only look at a few parameters, evaluate a few simple tests). In the future, could look at attestation for ECUs.
Select and tailor CCoA	5	CCoAs in the vehicle are purely automated; the vehicle operator is unaware of and indifferent to the cyber posture of vehicle systems. CCoAs in the vehicle involve turning off non-critical systems.	Baseline: 2 CANStomper: 3	The range of CCoAs in vehicles is currently limited to a few systems beyond "limp home" mode. For example, if traction control is assessed as damaged, it can be automatically disabled; if power steering fails, a fail-safe mode for steering engages. CANStomper removes messages from unidentified or illegitimate sources from the CAN bus.
Fail over to replicated resources	0	Not currently implemented (except for manual failover to physically separate handbrake). Could be part of a future vehicle, but would be very expensive.	0	
Switch communications to use alternative communications paths (e.g., different protocols, different communications media)	3	Switch from one CAN bus to an alternate to eliminate or reduce the effects of a bad actor on the first CAN bus. (If one ECU denies service on one bus but not another, switching can be helpful.)	Baseline: 0	In the vehicle under consideration, have two high-speed CAN buses and one low-speed CAN bus, but no switching is performed.

Locate and switch over to alternative mission data sources	0	Switch between different sensors (e.g., radar, lidar) to assess location - capability is critical to purely automated vehicles, but not applicable to this use case.	0	N/A
Locate and switch over to alternative information stores	0	N/A	0	N/A
Ensure that ongoing functioning is correct	*Sub-Objective Priority Rating:* 4	*Restatement / Rationale for Priority Rating:* Ensure that critical vehicle systems are functioning within specifications and characteristic behavior.	*Sub-Objective Score:* Baseline: 11 CANStomper: 37 Refresh: 11 Attestation: 26	
Activity	*Activity Priority Rating*	*Restatement / Rationale for Priority Rating*	*Activity Performance Score*	*Rationale for Performance Score*
Validate provenance of mission-critical and system control data	5	Validate the provenance of mission-critical and system control data from within the vehicle.	Baseline: 0 CANStomper: 5	Not done in vehicle under consideration. CANStomper validates the source of messages.
Validate data integrity / quality to ensure it has not been corrupted	5	Validate messages on CAN bus to ensure that each message is correct, does not demand impossible behavior.	Baseline: 2	Have maximum values for speed, RPM.
Validate software / service integrity / behavior to ensure it has not been corrupted	5	Validate ECU OS is correct and not corrupted.	Baseline: 0 Attestation: 3	Not done in vehicle under consideration. Attestation validates the integrity of the infotainment system software.
Validate hardware / system integrity / behavior to ensure it has not been corrupted	4	Validate that hardware is behaving within normal parameters.	Baseline: 0	Tire pressure monitoring system (TPMS) validates physical component, but not ECU hardware.
Constrain				
Identify potential damage	*Sub-Objective Priority Rating:* 5	*Restatement / Rationale for Priority Rating:* Identify potential damage to, and circumstances which could result in harm to, critical systems.	*Sub-Objective Score:* Baseline: 34 CANStomper: 59 Attestation: 49	
Activity	*Activity Priority Rating*	*Restatement / Rationale for Priority Rating*	*Activity Performance Score*	*Rationale for Performance Score*

Identify potentially corrupted or falsified information	4	Determine whether CAN bus messages to critical systems are falsified or corrupted.	Baseline: 1 CANStomper: 4	Some manufacturers and third parties are implementing IDSs, but an IDS capability is not present in the vehicle under consideration. Error checking may be performed for adaptive cruise control. CANStomper identifies falsified messages.
Identify potentially compromised or faulty processes or services (i.e., those which can no longer be trusted)	5	Perform self-checking of correctness for critical services, i.e., running software (braking, engine, etc.).	Baseline: 1 Attestation: 3	Vehicle under consideration performs self-check of all critical services upon turning on the vehicle. The quality of the check is unknown. In the future, could use stimulus-response checking. Attestation of infotainment system software determines whether that system is no longer trustworthy.
Identify potentially faulty, corrupted, or subverted components	5	Perform self-checking of correctness for critical components, i.e., hardware.	Baseline: 3 CANStomper: 4	Vehicle under consideration performs self-check of all critical components upon turning on the vehicle. (Checks brake fluid pressure, oil pressure, engine running, battery status, tire pressure.) In the future, could use attestation. CANStomper identifies when an ECU starts chattering with invalid data, as well as when a subverted component sends messages it is not authorized to send.
Isolate resources to limit future or further damage	*Sub-Objective Priority Rating*: 5	*Restatement / Rationale for Priority Rating:* Isolate critical systems from malfunctioning or harmful non-essential systems.	*Sub-Objective Score:* Baseline: 0 CANStomper: 20	
Activity	*Activity Priority Rating*	*Restatement / Rationale for Priority Rating*	*Activity Performance Score*	*Rationale for Performance Score*
Isolate an enclave or set of cyber resources suspected of being compromised or in a faulty state (e.g., to contain adversary activities, to prevent use of suspect information)	5	Isolate non-essential systems (in particular, infotainment) from critical systems to prevent them from causing harm.	Baseline: 0 CANStomper: 2	No capability in vehicle under consideration. CANStomper provides limited virtual isolation (quarantine).

Isolate a critical or sensitive enclave or set of cyber resources to defend against potential compromise, faults, or failures from other resources	5	Isolate critical systems from malfunctioning or harmful non-essential systems to protect them from harm.	Baseline: 0	No capability in vehicle under consideration.

Appendix B Abbreviations and Acronyms

ATT&CK™	Adversarial Tactics, Techniques & Common Knowledge
C2	Command and Control
CAN	Controller Access Network
CCoA	Cyber Course of Action
CERT	Computer Emergency Response Team
CIS	Center for Internet Security
CJA	Crown Jewels Analysis
CNSS	Committee on National Security Systems
CNSSI	CNSS Instruction
CoA	Course of Action
CONOPS	Concept of Operations
COOP	Continuity of Operations
COTS	Commercial Off-The-Shelf
CPS	Cyber-Physical System
CR	Cyber Resiliency
CRDP	Cyber Resiliency Design Principles
CREF	Cyber Resiliency Engineering Framework
CRM	Customer Relationship Management
CSF	[NIST] Cybersecurity Framework
CSG	Cyber Security Game
CSIAC	Cyber Security and Information Systems Information Analysis Center
CTF	Cyber Threat Framework
DBMS	Database Management System
DoD	Department of Defense
DoS	Denial-of-Service
ECU	Embedded Control Unit
EIT	Enterprise IT
GPS	Global Positioning System
I&W	Indications and Warnings
ICS	Industrial Control System
IdAM	Identity and Access Management
IDS	Intrusion Detection System
IEEE	Institute of Electrical and Electronics Engineers

IPS	Intrusion Protection System
ISO	International Standards Organization
IT	Information Technology
KES	Keyless Entry System
KPP	Key Performance Parameter
KSA	Key System Attribute
LSPE	Large-Scale Processing Environment
M&S	Modeling and Simulation
MBE	Model-Based Engineering
MBSE	Model-Based Systems Engineering
MECR	Measuring the Effectiveness of Cyber Resiliency
MIA	Mission Impact Analysis
MIP	MITRE Innovation Program
MOE	Measure of Effectiveness
MOP	Measure of Performance
N/A	Not Applicable
NIST	National Institute of Standards and Technology
NSA/CSS	National Security Agency / Central Security Service
NTCTF	NSA/CSS Technical Cyber Threat Framework
OBD	On-Board Diagnostics
ODNI	Office of the Director of National Intelligence
OEM	Original Equipment Manufacturer
PII	Personally Identifiable Information
PIT	Platform IT
RPM	Revolutions Per Minute
SAE	Society of Automotive Engineers
SDLC	System Development Lifecycle
SME	Subject Matter Expert
SOC	Security Operations Center
SP	[NIST] Special Publication
SoS	System-of-Systems
SSM-CR	Situated Scoring Methodology for Cyber Resiliency
TPMS	Tire Pressure Monitoring System
TTP	Tactic, Technique, or Procedure
TTPs	Tactics, Techniques, and Procedures
UI	User Interface

USB	Universal Serial Bus