TLP:CLEAR

SECURE CLOUD BUSINESS APPLICATIONS (SCUBA)

EXTENSIBLE VISIBILITY REFERENCE FRAMEWORK

June 2023

Cybersecurity and Infrastructure Security Agency

TLP:CLEAR

Contents

Executive Summary

Executive Order 14028, *Improving the Nation's Cybersecurity*, seeks "to improve [the Federal Government] efforts to identify, deter, protect against, detect, and respond to [malicious] actions and actors." To achieve its mission and strengthen cybersecurity across the Federal Government, the Department of Homeland Security (DHS) Cybersecurity and Infrastructure Security Agency (CISA) requires visibility across various Federal Civilian Executive Branch (FCEB) agency domains. This visibility enables CISA to develop and share insights across the FCEB, ensuring that CISA can identify threats; protect against potential attacks; and perform hunt, incident response, and analysis activities.

The Office of Management and Budget (OMB) released M-21-31 in accordance with and to address the requirements in Section 8 of Executive Order 14028. M-21-31 includes requirements for CISA to provide guidance to agencies in developing a schema for sharing their logs and to publish tools to help agencies facilitate their assessment of logging maturity. The extensible Visibility Reference Framework (eVRF) supports these requirements by providing a framework for organizations to identify visibility data that can be used to mitigate threats, understand the extent to which specific products and services provide that visibility data, and identify potential visibility gaps. Agencies can then use this knowledge to assess their logging maturity, articulate their visibility, direct resources to close visibility gaps, and enhance overall visibility into potential threats.

eVRF Layout

The eVRF is divided into the *eVRF Guidebook* (this document) and eVRF workbooks. The Guidebook is an instruction manual for eVRF; it defines and describes key concepts, roles and responsibilities, and workflows. Each eVRF workbook defines a visibility surface and enables organizations to produce their own visibility coverage maps for as-planned or as-implemented system configurations. Additionally, organizations can use coverage maps to identify desired visibility or visibility requirements. Figure 1 presents the eVRF document structure.

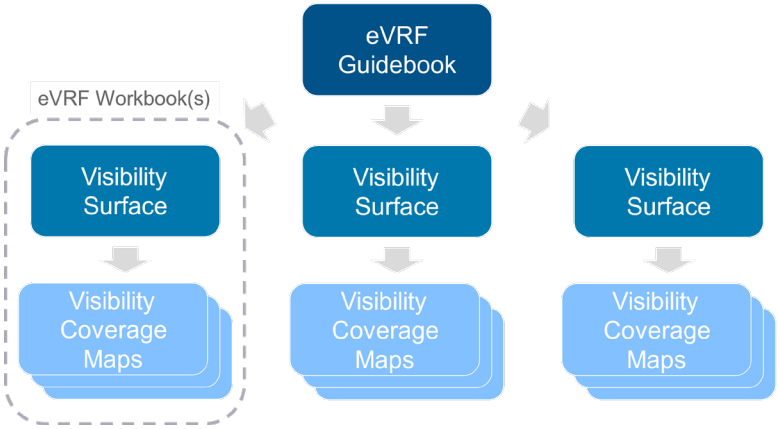

Figure 1: eVRF Document Structure

As organizations work through the workbooks, visibility coverage maps will be populated. Organizations can combine the completed coverage maps into visibility coverage comparisons. These comparisons provide a quick visual reference that can help identify where coverage gaps might exist. Visibility coverage comparisons can also be created to allow organizations to analyze and gain insights into their visibility across their enterprise.

The eVRF was developed for organizations to identify and evaluate visibility in digital environments. CISA will use this framework to communicate telemetry requirements with Federal Civilian Executive Branch Agencies.

1. Introduction

1.1 eVRF OVERVIEW

The Cybersecurity and Infrastructure Security Agency (CISA) requires visibility across various Federal Civilian Executive Branch (FCEB) agency domains. This visibility enables CISA to develop insights to share across the FCEB, ensuring that CISA can identify threats; protect against potential attacks; and perform hunt, incident response, and analysis activities.

eVRF Purpose

The purpose of the extensible Visibility Reference Framework (eVRF) is to provide a framework for organizations to identify visibility data that can be used to identify and mitigate threats, understand the extent to which specific products and services provide that visibility data, and identify potential visibility gaps.

eVRF Goals

The eVRF has the following five goals:

Goal 1: Communicate requirements for FCEB agencies to provide CISA with the necessary data to protect agency networks, devices, cloud-based environments, data, and systems.
Goal 2: Enable agencies to (a) evaluate their ability to collect relevant visibility data and (b) model their coverage of CISA's visibility requirements.
Goal 3: Promote partners' ability to incorporate key visibility concepts into their own threat-informed cyber practices and security baseline configurations.
Goal 4: Provide a framework for organizations to evaluate the visibility of products' capabilities and features and to characterize the visibility gaps that various products can fill.
Goal 5: Meet OMB M-21-31 requirements for CISA to provide guidance to agencies in developing a schema for sharing their logs and to develop and publish tools that help agencies facilitate their assessment of logging maturity.

1.2 BENEFITS OF EVRF

The eVRF provides the following benefits to organizations that adopt the framework:

1. Provides a model to characterize visibility across a broad set of domains that is representative of an organization's modern enterprise.
2. Informs an organization's situational awareness and enables organizations to prioritize the collection and analysis of visibility data across their enterprises to best mitigate against threats and improve their overall risk posture.
3. Allows for the identification of gaps in visibility coverage and enables the establishment of new targets and/or system configurations capable of addressing visibility needs.
4. Informs procurement decisions by providing perspective on visibility and impact prior to implementing a product and/or its system baseline configuration settings.
5. Provides a dynamic methodology, which can include new domains and telemetry as ecosystems continue to evolve.

1.3 DOCUMENT ORGANIZATION

The *eVRF Guidebook* (this document) is part of the larger eVRF document set. This set is considered a library, which will continue to grow over time as new domains are identified. The Guidebook identifies visibility as a unique characteristic of cybersecurity and includes a structure and workflow to characterize visibility for different portions of a cyber system. The Guidebook is an instruction manual that informs users about eVRF's key concepts, roles and responsibilities, and workflows.

This document is organized into the following five sections and three appendices:

- Section 1 provides basic scoping information that articulates the intention and focus of the document.
- Section 2 discusses the key concepts of visibility that informed the creation of the eVRF.
- Section 3 provides generalized guidance regarding how to apply an eVRF workbook.
- Section 4 explains how agencies and CISA will apply an eVRF workbook.
- Section 5 provides conclusions.
- Appendix A discusses how the eVRF relates to other CISA programs.
- Appendix B defines key terms used throughout the document.
- Appendix C discusses key background documents.

1.4 INTENDED AUDIENCE

The *eVRF Guidebook* is designed by CISA to define concepts, requirements, and mechanisms for collecting, evaluating, and analyzing telemetry for communication with federal civilian agencies, service providers, and other public and private sector partners. Organizations may also leverage the *eVRF Guidebook* as reference for their analysts, solution architects, and cybersecurity acquisition decision-makers to make threat-informed decisions on visibility and improve their ability to hunt for threats and investigate incidents across their enterprise. Organizations can use this document to evaluate technology solutions (including both open-source and for-profit vendors) to express the visibility that such products offer and to identify the product tiers, add-on capabilities, and configuration settings needed to meet CISA's requirements. Even though this framework was developed by CISA for CISA stakeholders, any organization interested in incorporating visibility into their cybersecurity practices or communicating visibility requirements and gaps can use the concepts and workflows in eVRF.

1.5 ASSUMPTIONS AND CONSTRAINTS

This Guidebook describes the concepts, processes, and scope of eVRF. Individual eVRF workbooks, produced on a case-by-case basis, will describe specific visibility requirements. Currently, this Guidebook recognizes that as-built agency systems may not fully align with visibility requirements, but that agencies will satisfy the various roles and responsibilities of eVRF over time. Full eVRF implementation may require updates to products, services, or service level agreements, as well as additional expertise or training. Agencies will need to work with their solution providers and CISA as service and product providers evolve and extend their services and capabilities to accommodate customers' visibility needs.

This Guidebook does not constitute a request for product proposals or solicitations; nor should this Guidebook be seen as detailed specifications or formal requirements for vendors or service providers. The terms and details of eVRF are subject to change at any time. Furthermore, this Guidebook does not supplant or supersede any previously issued CISA guidance, government-wide policies, or applicable law. Agencies should continue to comply with telemetry and logging requirements, including those that require agencies to provide network visibility or allow agencies to provide cloud telemetry. Agencies remain the sole data owners for all telemetry data that they generate; agencies are merely sharing visibility of that data with CISA. CISA will use eVRF in conjunction with the MITRE ATT&CK© Framework[1] to develop specific threat models and methodologies; the MITRE ATT&CK© Framework is developed and maintained outside the scope of eVRF activities. Agencies can use eVRF to characterize the visibility and completeness of observation coverage but realizing the full benefit of eVRF depends on employing other systems, methods, and platforms for risk-based countermeasures, determining the efficacy of mitigations, and collecting/processing the sensor data to derive value from the observations.

1.6 RELATIONSHIP TO OMB M-21-31

All organizations can leverage the eVRF to assess and enhance their organizational visibility. However, within the FCEB, the eVRF can also help agencies comply with OMB M-21-31.

OMB released memorandum M-21-31[2] ("*Improving the Federal Government's Investigative and Remediation Capabilities Related to Cybersecurity Incidents*") on logging, log retention, and log management for FCEB Agencies in support of the Executive Order 14028 *Improving the Nation's Cybersecurity*.[3] The memo includes a maturity model for event log management and logging requirements for many log categories across an enterprise.

M-21-31 includes a series of requirements related to CISA. These include requirements to provide guidance to agencies in developing a schema for sharing their logs and to develop and publish tools to help agencies facilitate their assessment of logging maturity across their organization. The eVRF and its associated work products will provide guidance to agencies on developing schemas for sharing logs from various log categories. The eVRF can also help agencies assess their logging maturity by understanding the visibility provided by their logs and identifying visibility gaps.

M-21-31 also requires agencies to provide relevant logs to CISA upon request and states that real-time access to that data may be required. The eVRF and its associated work products can help facilitate those transactions by ensuring required visibility is identified and prioritized.

---

[1]"MITRE ATT&CK," MITRE ATT&CK, 2015–2022, https://attack.mitre.org/

[2]Executive Office of the President Office of Management and Budget, August 27, 2021, https://www.whitehouse.gov/wp-content/uploads/2021/08/M-21-31-Improving-the-Federal-Governments-Investigative-and-Remediation-Capabilities-Related-to-Cybersecurity-Incidents.pdf, M-21-31,.

[3]Executive Order 14028, "*Improving the Nation's Cybersecurity*", (May 2021). https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/.

1.7 EVRF USE WITHIN FEDERAL ACQUISITION LIFECYCLES

When federal agencies apply the eVRF, they should consider the acquisition, procurement, and contract implications. Agencies can use eVRF during the acquisition lifecycle to evaluate products or services that generate cybersecurity telemetry data (e.g., logs). Since the relationships between agencies and vendors are governed by procurements, this Guidebook does not obligate or otherwise require any vendor to use eVRF. Any requirement for a vendor to use eVRF techniques or mechanisms (e.g., produce and update product visibility coverage maps, submit information for eVRF workbooks, validate information with CISA, or any other eVRF activity undertaken on behalf of any agency) must be set out in one or more agency contracts, potentially requiring modifications to established contracts.

Agencies could use the eVRF as a part of market research early in the acquisition lifecycle and can use publicly available materials to perform market research analysis. Vendors may even find it advantageous to provide eVRF coverage maps as it provides an avenue of communicating the capabilities of their products and services in a way that makes sense to federal entities. Incorporating eVRF into contract requirements and solicitations would allow an organization to identify its own visibility gaps and enable informed decision-making when selecting a vendor.

Within existing contracts, agencies may be able to coordinate service configuration with their current vendors to enable generation and delivery of the appropriate security event and telemetry data. Agencies may also use eVRF when issuing new contracts—for example, when procuring capabilities or services to address any identified risks. Insights from the eVRF process can inform the design of an agency security architecture that combines products and services from multiple vendors to provide comprehensive solutions.

2. Visibility

This section defines and introduces key concepts to ensure users have a common understanding of the eVRF. To promote understanding, this section uses a scenario of a high-value physical asset to draw parallels with cyber-systems and assets.

2.1 KEY VISIBILITY CONCEPTS

Visibility

In its most general sense, the term *visibility* describes something that is visible. CISA applies the term visibility to refer to (a) the observable artifacts of digital events and (b) the characteristics of the digital environment in which those events take place. By collecting and analyzing the observable artifacts and characteristics of an environment, organizations will have the data necessary to conduct forensic investigations into threat activity and maintain better awareness of activity on an ongoing basis. Desired qualities for visibility data include the cost-effective and scalable collection of relevant data, the ability to receive data at cyber-relevant speeds, etc. The more in-depth and extensive the technical visibility, the greater opportunity for an organization to detect high-priority threats to networks, devices, and data.

Visibility provides context-specific insights about the activity taking place within a given environment. Because it is context-specific, the types of data that provide visibility will vary across the enterprise. For instance, within a cloud-centric context, the most useful visibility may come from cloud API activity logs, but biometric event logs may be preferable for visibility into mobile device behaviors. The heterogeneity of data types across contexts can make obtaining consistent visibility across an entire enterprise difficult. The eVRF suggests dividing the enterprise into multiple visibility surfaces, each centered around a different type of system with its own context. The next section discusses visibility surfaces in more detail.

Many visibility mechanisms have already been deployed across enterprise architectures through implementation of security controls associated with standards such as NIST SP 800-53 *Security and Privacy Controls for Information Systems and Organizations*. While implementing eVRF across all domains within an organization's enterprise can be a lengthy effort, capturing the visibility associated with existing security control requirements can improve familiarity with the workflow and increase efficiency for subsequent analysis. This would also enable an organization to begin documenting and understanding the visibility that currently exists and where they may focus initial efforts to identify gaps in visibility.

Visibility refers to the observable artifacts of digital events and the characteristics of the digital environment in which those events take place. Visibility provides context-specific insights about activity within a given environment.

Visibility Surface

A *visibility surface* refers to a digital environment where cyber-observable data exists or should exist and is therefore an environment-specific instantiation of visibility. In the same way that an attack surface is comprised of many different points from which a system can be attacked, a visibility surface is made up of many observation points, or perspectives, from which a system can be observed. As detailed in the section below, observation points provide architectural context, which tightly couples visibility surfaces to real data common to the domain. The cyber-observable data—logs, configuration settings, packet data, etc.—that contribute to a visibility surface are essential for providing evidence of malicious activity.

Figure 2 displays the visibility surface of the target system through highlighting the space within the fenced in area in orange. The high-value assets are serving their business need and meet their design intent. An understanding of malicious actors and the anticipated approaches they would employ to gain access to the asset can be derived from the value of the asset and the context of its placement.

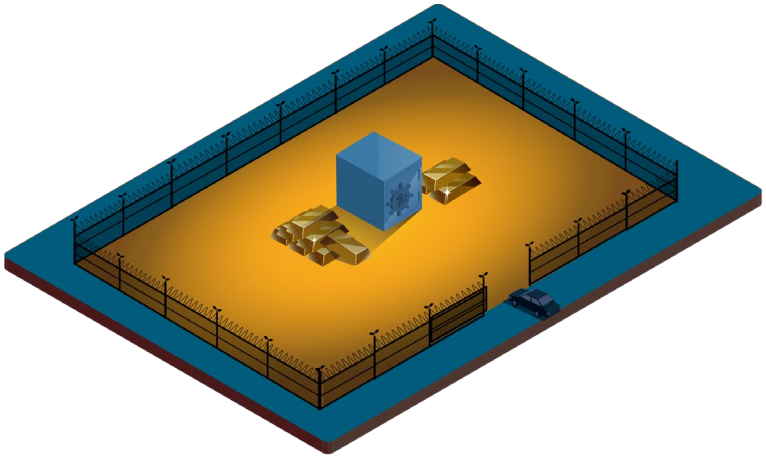

Figure 2: Visibility Surface—Environment for Where Data Exists or Should Exist

Within eVRF, visibility surfaces allow an organization to identify which data can be used to recognize threat actor tactics, techniques, and procedures (TTPs) within a system. In addition to identifying relevant data and TTPs, each eVRF visibility surface is scoped to a particular type of digital environment (e.g., cloud business applications, workstation operating systems). Once these parameters of a visibility surface are defined (see Section 3.1), organizations can overlay additional information to produce coverage maps that portray the visibility provided by one or more system configurations.

Observation Point

An *observation point* defines the architecture location of a telemetry source in the given domain. For example, the following are all possible observation points in a cloud architecture: the Cloud Service Provider (CSP), the Cloud Access Security Broker (CASB), any Security-as-a-Service (SECaaS) solution, and virtual network locations throughout. An observation point may be the sensor positioning within a cloud or network topology or a specific host for endpoint visibility. An observation point can be in the same architecture location that policies are applied and is often associated with a policy enforcement point (PEP) and/or a policy decision point (PDP). Observation points may be in line with data, at data entry, or at data exit for a domain. Collecting telemetry from multiple observation points increases the breadth of visibility across a domain.

> In Figure 3, the concept of observation points is represented by guard towers in each corner of the fenced area. The observation points can host one or more sensors, which provide visibility into the visibility surface. The location of the observation point impacts the visibility available to sensors hosted at that location.

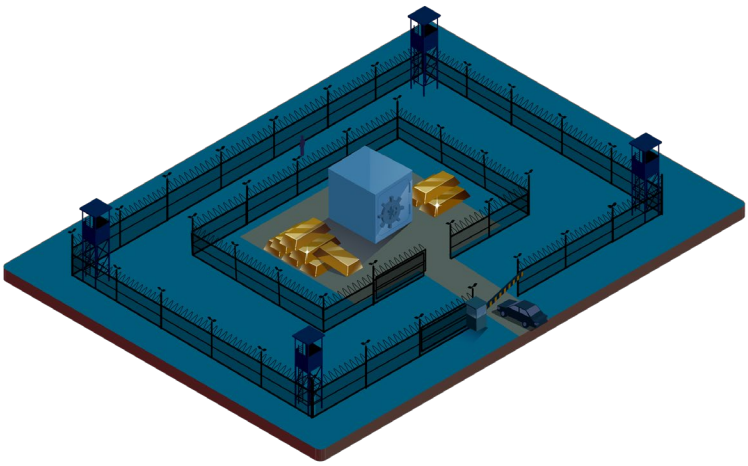

Figure 3: Observation Points—Architecture Locations for Telemetry Sources.

Sensors

Sensors collect telemetry at observation points. Multiple sensors may be co-located at the same observation point. Sensors should be selected and deployed to provide specific insights. When they share an observation point, they ideally produce complementary data, which augment and enrich each other. For example, an organization may have both a Web Application Firewall (WAF) and a Next Generation Firewall at the same observation point (gateway), and both firewalls together may provide greater insight into network activity.

> The various light sources shown in Figure 4 display different sensors that each observation point provides. Additional sensors can increase the amount and type of visibility an organization has on the asset, as well as provide an increase of visibility detail when coverage is overlapped, such as the figure's drone mounted purple sensor overlapping with the top-right tower's light blue sensor.

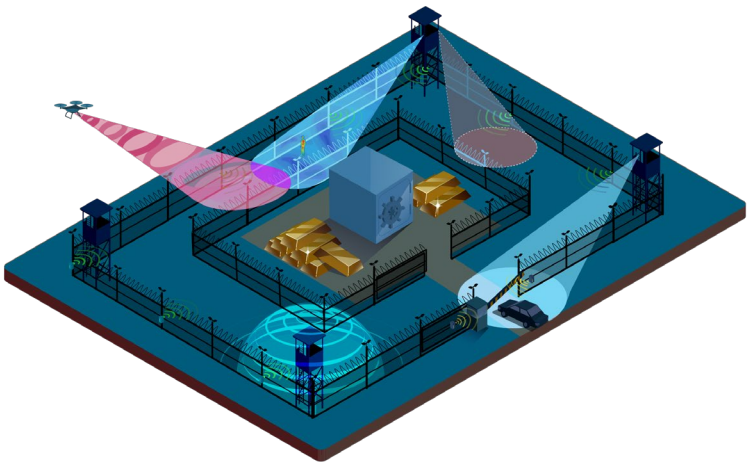

Figure 4: Sensors—Positioned at Observation Points and Provide Telemetry

Visibility Coverage Maps

A *visibility coverage map* characterizes the ability of a product or organization to address a visibility surface by providing relevant cyber-observable data. Whereas a visibility surface describes the scope of the environment and its relevant data and TTPs, a visibility coverage map conveys the extent to which available data provides sufficient visibility into cyber threat activity.

Using eVRF, organizations create coverage maps by using an eVRF workbook to indicate the data currently or potentially available in the environment. Organizations can create coverage maps for each actual or presumed tiered logging level of a vendor's major product offering. Organizations should periodically update coverage maps to accurately describe rapidly changing telemetry options. The workbook will use that input to produce a color-coded visualization that shows which MITRE ATT&CK techniques are addressed by the available data. Organizations can also show metrics to indicate the quality of coverage for each technique. Bolstering an organization's coverage map in a visibility surface builds crucial security event context for detection and mitigation.

> The purple light source shown in Figure 5 displays the coverage of the visibility surface provided by a sensor at a single observation point. The product coverage map can include in its description the limitations of the sensor; ideal usage characteristics; and any licensing details, sensor upgrade, or complimentary enrichment options.

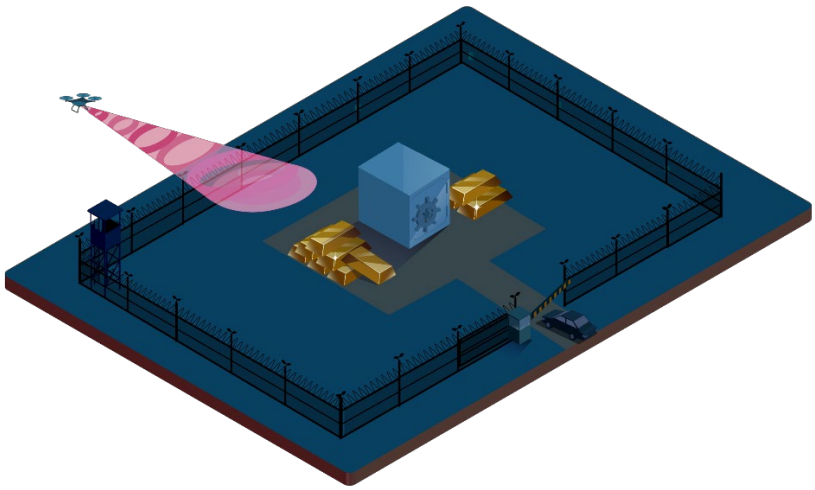

Figure 5: Product Coverage Map—Visibility From a Single Observation Point and Sensor Type

Visibility Requirements Maps

A *visibility requirements map* is a special purpose coverage map used for the identification of cyber-observable data, which must be shared between parties for common situational awareness and use (for a given visibility surface). The visibility requirements map can identify the criticality of the sharing for given metadata, the diversity of observation points and sensor inputs required, or other "cyber-observable data quality" attributes.

In Figure 6, the orange coloring throughout the fenced area represents the visibility requirements. The requirements set is agnostic of the observation points and sensors offered by any given vendor, but instead can focus on the visibility surface—the use of the high-value asset and its environment.

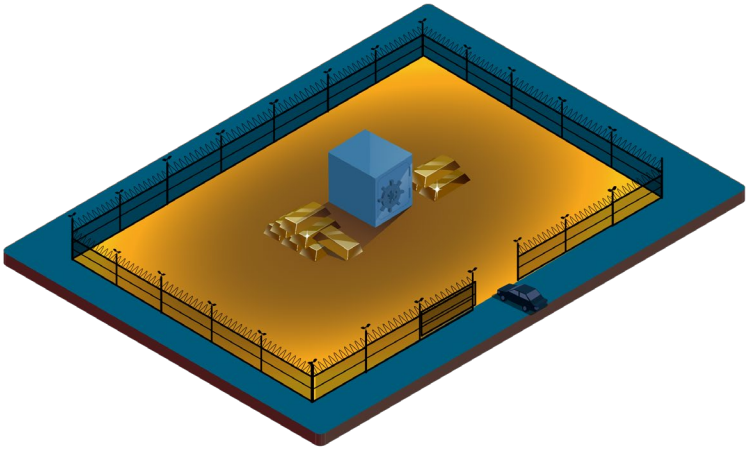

Figure 6: Visibility Requirements Coverage Map—Data Sharing for Common Situational Awareness

By using visibility requirements maps, an authoritative organization (e.g., CISA) can communicate telemetry requirements to other participating organizations for a given visibility surface while staying agnostic to any vendor's implementation. The organization should update their requirements maps as their understanding of threats changes, as visibility capabilities within the domain evolve, and as other organizations mature in their own telemetry use (and rely less on the authoritative organization's supplemental protections).

Visibility Coverage Comparisons

A *visibility coverage comparison* consists of two or more coverage maps overlaid onto the MITRE ATT&CK framework for evaluation of competing and/or complementary products and services. The coverage comparison map answers the question, "For which ATT&CK technique does the combination of products/services produce telemetry data?"

Visibility coverage comparisons can be treated as nominal stand-ins for organizations' as-built technologies or for proposed architectures being considered for deployment. Organizations can create a visibility coverage comparison when considering competing products or multiple security architectures; a visibility coverage comparison can show a side-by-side comparison of available telemetry data in each product offering or architecture design. Visibility coverage comparisons are tools for determining the prioritization of telemetry and return on investment in telemetry options. They are the culmination of an eVRF workbook analysis and are used to produce high-level insights.

2.2 DIVISION OF ENTERPRISE INTO DOMAINS

An organization's digital enterprise is extensive; it includes many different hardware devices, networks, virtual environments, operating systems, and applications. Prior to defining an eVRF visibility surface in a workbook, it is helpful to segment the enterprise into components to create a manageable scope for a single visibility surface. There are many ways to categorize or scope visibility surfaces within an enterprise, resulting in different approaches for each organization, as they may need to account for specific characteristics, structures, or other considerations.

Segmenting an enterprise into components will assist in creating a comprehensive list of the existing observation points and their related sensors. First, it is helpful to visualize the relationship of domains, observation points, and sensors in a hierarchical view, as seen in the example in Figure 7. The sensors are grouped within the observation points, which are aligned to domains.

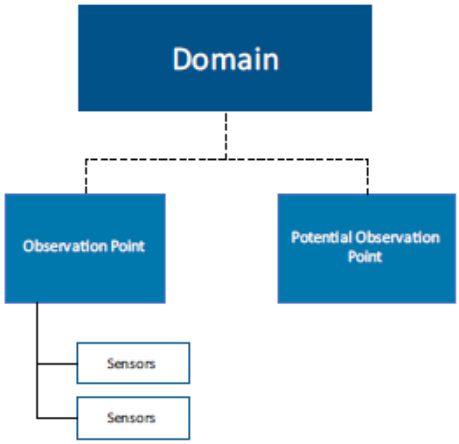

Figure 7: Domain Hierarchy

An organization should choose a set of common attributes to scope to a domain with an assigned title. Since a domain is a collection of shared observation points (architectural locations for sensors), this is often tightly coupled with the deployment model in use. Hence, the network topology, data forwarding path, and service delivery model can all help to inform the scope (whether any given observation point is in or out) of a domain. Figure 8 illustrates how "remote work" and "email" may be used for identifying a group of observation points within an organization.

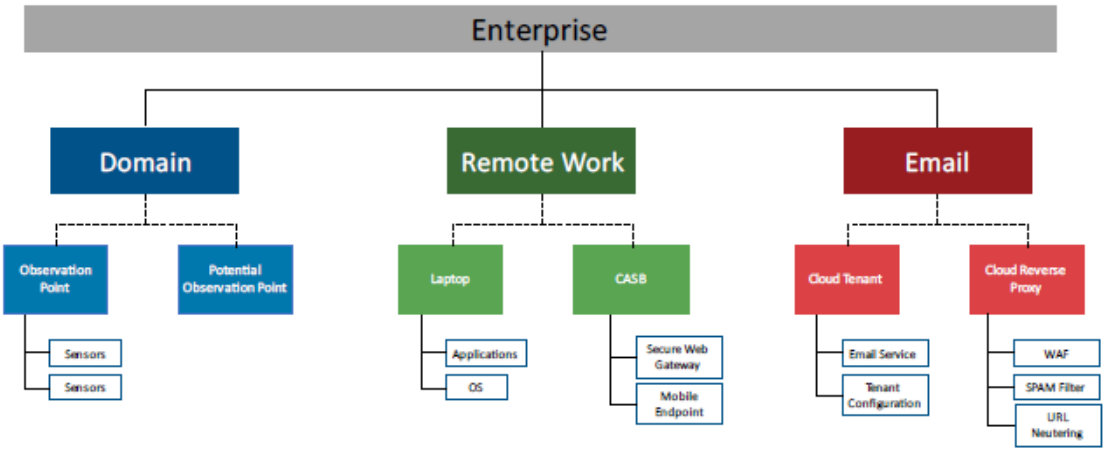

Figure 8 Components of a Remote Work and Email Domain

After an enterprise has been divided into domains, the process of creating a visibility surface is simplified, as the observation points and sensors have already been identified and grouped. There are many possible common attributes that could be used to scope a set of domains into visibility surfaces. Some examples include:

- Physical location
- Organization structure

- Administrative domain
- Shared need-to-know
- Existing logical grouping
- Any combination of the above

A visibility surface may or may not incorporate each of the domain's identified observation points, as seen in Figure 9. Both the scope and level of detail of a visibility surface will determine the extent of its intersection with a given domain. Similar to how an enterprise can be segmented into multiple domains to manage complexity, leveraging multiple visibility surfaces can assist in managing the workload for eVRF visibility characterization.

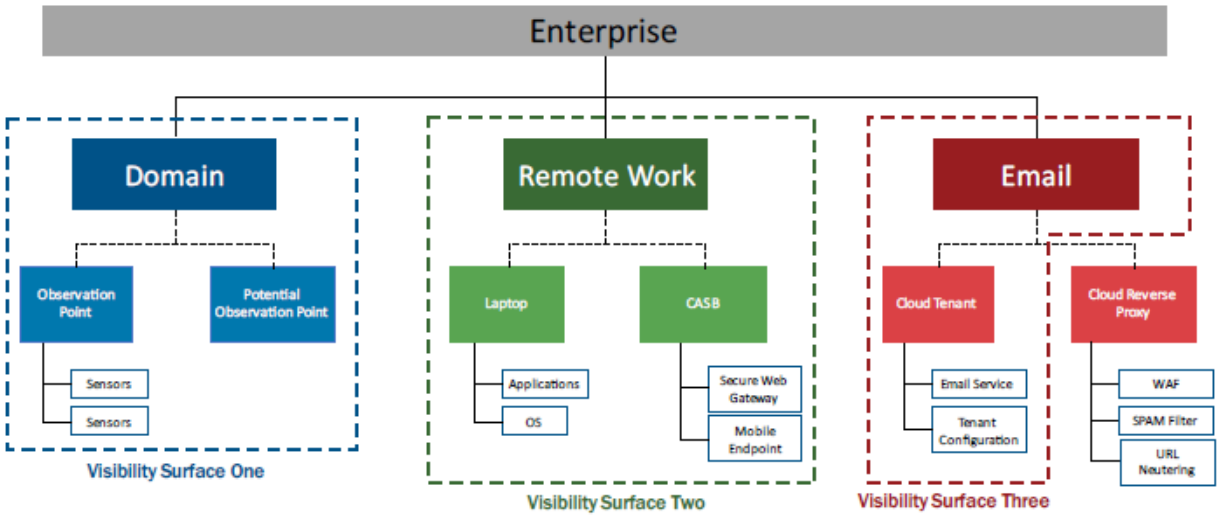

Figure 8: Visibility Surface Scopes

In this example, the "Email" domain consists of candidate observation points for the "Cloud Tenant" as well as "Cloud Reverse Proxy" architecture locations. These architectures can be evaluated for inclusion within the visibility surface scope and included or excluded based on desired common attributes. In this case, the visibility surface scope sought to maximize reuse by multiple organizations and limit the coverage to a single service provider. As a result, visibility surface three includes the observation point for "Cloud Tenant," but determined the "Cloud Reverse Proxy" would not always be present and therefore excluded it from the characterization.

3. Generating an eVRF Workbook

An eVRF workbook defines specific visibility surfaces and enables organizations to produce their own visibility coverage maps for as-planned or as-implemented system configurations. By using the workbook to identify what visibility data is available in their environment, organizations can identify visibility gaps and set visibility requirements. Vendors may also provide product-specific visibility coverage maps to indicate the visibility offered by individual products or product tiers.

An eVRF workbook offers a flexible way to create and edit a visibility surface definition and coverage maps. An interactive workbook application is currently in development.

As organizations develop each workbook, visibility coverage maps will be populated within the workbook. These maps will provide a quick visual reference showing potential gaps in coverage.

Figure 10 shows how several types of visibility coverage maps are developed for each visibility surface and how each layer provides unique insights.[4] The color-coding provides a visual reference of how well each MITRE ATT&CK technique is addressed within the workbook.

- **CISA Visibility Requirements Map**: The CISA visibility requirements coverage map is developed by CISA to show telemetry generation, collection, and processing requirements.
- **Product Coverage Maps**: Product coverage maps can be developed by vendors or service providers to show how the visibility of their solutions informs ATT&CK TTPs.
- **Environment Coverage Maps**: Environment coverage maps characterize the organization's as-built or to-be-built environments and may be produced using one or more product coverage map(s). Environment coverage maps consider factors like product configuration and licensing level to accurately reflect the visibility provided by the organization's implementation of visibility products.

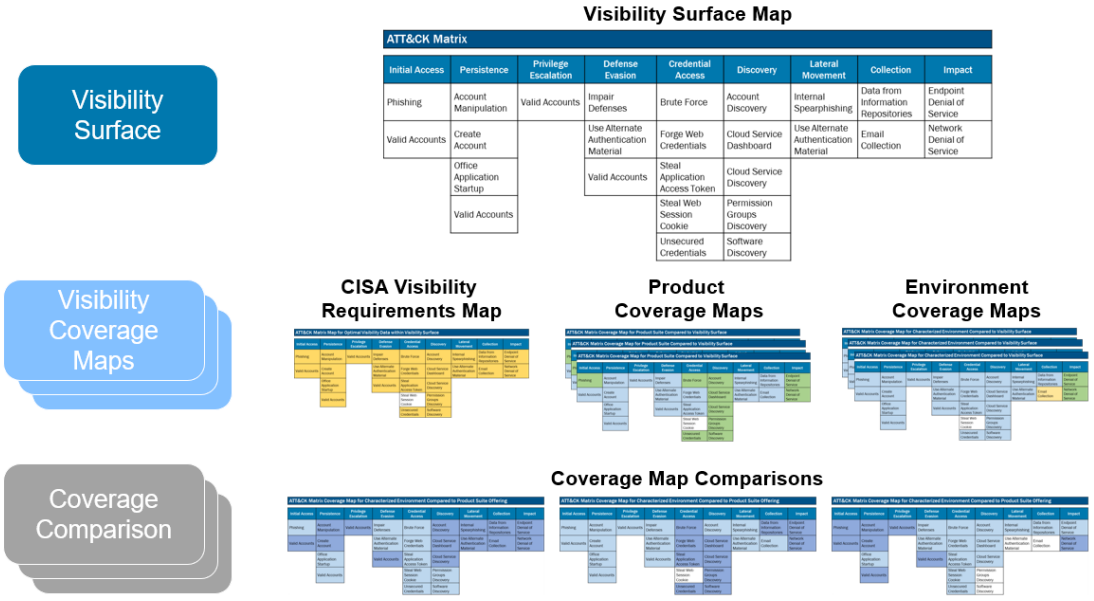

Figure 9: eVRF Workbook Structure

[4] The visibility surface map used in this figure is one example of how a visibility surface map can be derived from MITRE ATT&CK to represent a specific domain. Different combinations of ATT&CK tactics and techniques will be used for different domains.

After deriving the coverage maps, an organization can generate a visibility coverage comparison by combining multiple coverage maps. The visibility coverage comparison can be used for analysis and to generate insights.

The eVRF workflow defined in this Guidebook refers to a complete workflow process. In practice, some visibility artifacts will exist and will not need to be recreated. Hence, as organizations employ this workflow and a library of artifacts grows, some of the steps may be bypassed.

3.1 WORKFLOW PROCESS OVERVIEW

The eVRF workflow describes the process for establishing a visibility surface and building coverage maps to evaluate the extent of visibility available in an environment. The process is separated into three phases:

- **Phase 1: Define Visibility Surface:** In this phase, a visibility surface definition is created, which establishes the surface boundaries and identifies the required visibility data. A visibility surface can be defined with one or more observation points containing one or more sensors each. Many organizations will choose to use an existing visibility surface definition instead of creating a custom or new definition.
- **Phase 2: Produce Visibility Coverage Maps:** In this phase, a coverage map is produced to characterize a selected environment to indicate whether available data provides the desired visibility. Some organizations may choose to produce multiple coverage maps to indicate varying levels of visibility in different parts of the environment. Many organizations will choose to develop coverage maps based on vendor-provided coverage information.
- **Phase 3: Generate Visibility Coverage Comparisons for Analysis and Insights:** In this phase, organizations analyze the coverage maps to identify gaps in coverage, to establish targets for new visibility data that must be collected, or to generate other operational or business insights. Organizations can produce and consolidate coverage comparison for multiple parts of the environment by combining coverage maps from more than one visibility surface.

Figure 11 shows the eVRF workflow. The following sections describe these phases in more detail.

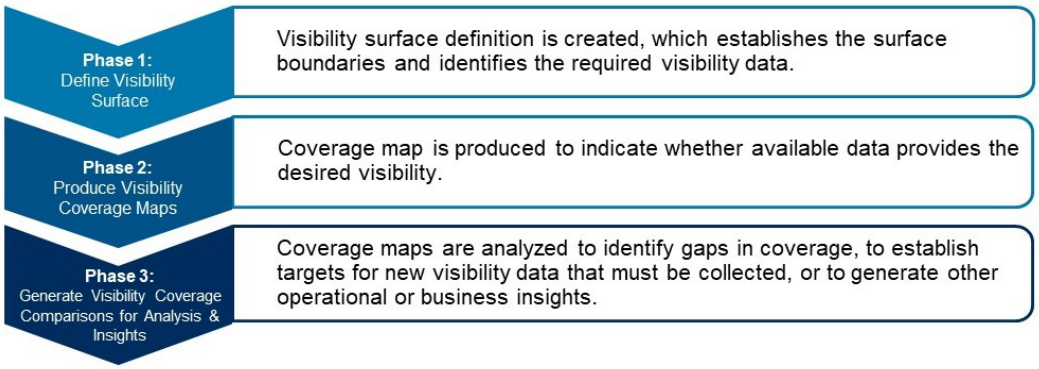

Figure 10: eVRF Workflow

Phase 1: Define a Visibility Surface

Organizations may choose to use an existing visibility surface definition, such as one published by CISA, or they may choose to create a new definition. As the set of established visibility surface definitions grows, the need for new definitions will be diminished. In practice, a visibility surface can be scoped with one or more observation points, containing one or more sensors each. Every eVRF workbook will need to define the visibility surface that will be examined in that workbook.

Each visibility surface definition identifies the following:

Scope: Identifies the bounds of the digital environment included in the visibility surface.
Relevant Data: Identifies which types of data are needed to provide evidence of threat actor TTPs.
ATT&CK Matrix: Identifies the ATT&CK techniques that are relevant for the environment.
ATT&CK-to-Data Overlay: Identifies which ATT&CK techniques are addressed by the relevant data types.
Create Templates for Coverage Maps: Prepares for subsequent phases by generating templates for data entry to characterize systems.

To use an existing visibility surface definition, locate the relevant eVRF workbook (e.g., cloud business applications) and skip to Phase 2. To create a new visibility surface definition, begin with a blank eVRF workbook template and conduct five sequential activities to fill required information into the workbook, as seen below in Figure 12.

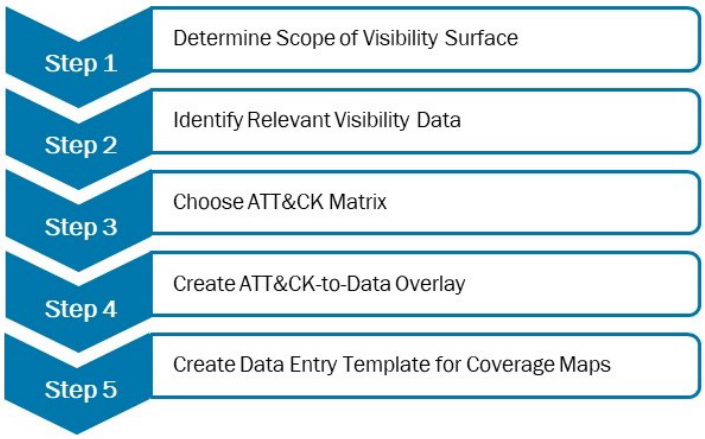

Figure 11: eVRF Workflow Phase 1

Phase 1, Step 1: Determine Scope of Visibility Surface

Establish the scope of the environment to be captured in the visibility surface definition. Consider both the type of environment (e.g., cloud business applications, endpoint detection and response capabilities) and the appropriate level of granularity within the technology stack (see Figure 13). Each increment down the technology stack provides an increased level of detail and greater precision when evaluating visibility. However, it also limits the scope of the visibility surface, reducing

opportunities for reuse and requiring additional visibility surfaces to be defined for full ecosystem awareness.

The scope of the visibility surface may limit the number of ATT&CK sub-techniques considered. Organizations should consider all ATT&CK sub-techniques when creating the visibility surface.

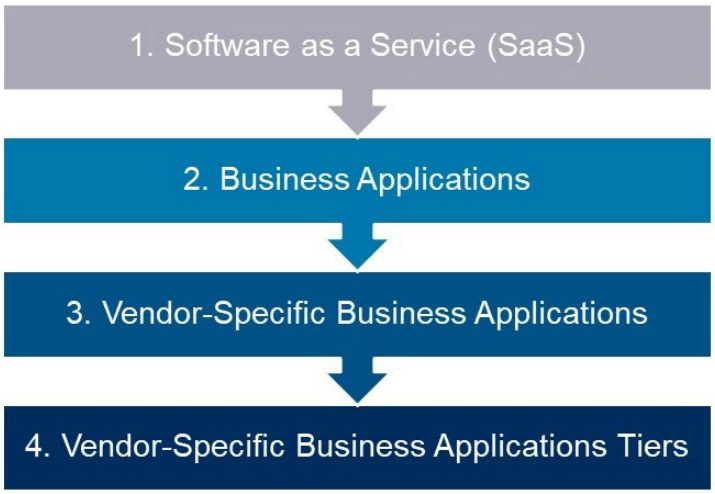

Figure 12: Visibility Surface Scoping Example

Phase 1, Step 2: Identify Relevant Visibility Data

Create a listing of the data that applies to the visibility surface. To produce an effective visibility surface definition, it is important that this activity identifies all the data desired for visibility into the technology domain. Organizations may need to engage several experts to participate in this activity to identify the necessary data. Include experts who have comprehensive experience with the technologies that are in scope as well as cybersecurity experts who can identify the types of data used to conduct forensic analysis of those technologies. Changes to the technology and threat environment may require changes to the list to preserve accuracy.

Organizations should organize the list of data into four sets with increasing levels of detail:

> **Category:** Identifies a component (i.e., application, software, service, etc.) of a system in which cyber-observable data exists (e.g., email, document management).
> **Event:** Identifies a process that occurs within the defined component (e.g., receive incoming email, sending outgoing email).
> **Metadata:** Lists individual data objects or information elements that document the state of the system, an event that occurred, and/or how it may have occurred (e.g., sender, recipient, subject).
> **Description:** Provides additional details or notes about the activity or data being logged.

Phase 1, Step 3: Choose ATT&CK Matrix

Determine the ATT&CK techniques relevant for the environment, potentially using a pre-defined MITRE ATT&CK Matrix (e.g., traditional or cloud).

Optionally, organizations may choose to increase the fidelity of their eVRF evaluation by conducting the evaluation at the level of "sub-technique" instead of "technique." If an organization chooses to evaluate sub-techniques, only the relevant sub-techniques need to be included. In this way, organizations can scale the fidelity of visibility assessments to accommodate each organization's needs and risk posture.

Phase 1, Step 4: Create ATT&CK-to-Data Overlay

Review the visibility data identified in Step 2 and determine whether the data can provide visibility into each of the ATT&CK techniques. Capture these assessments and use them to create an ATT&CK-to-Data overlay. As with Step 2, organizations may need to engage technology and cybersecurity experts to participate in this activity.

The completed overlay identifies the relevant visibility data for the visibility surface. Even without creating the visibility coverage maps and visibility coverage comparisons Phases 2 and 3 of the eVRF workflow describe, this overlay can provide valuable insight to guide decisions and awareness about how log data can be used to identify threat activity.

Phase 1, Step 5: Create Data Entry Template for Coverage Maps

Create the templates for Phase 2, which organizations will use to characterize their environment and identify the available visibility data.

In Phase 2, organizations will use this data entry table to indicate whether each service or application in their environment provides the desired visibility data. The resulting information will be displayed as a visibility coverage map.

Phase 2: Produce Visibility Coverage Maps

Visibility coverage maps enable organizations to analyze and communicate information about the visibility the data provides in a given environment. Organizations can use an eVRF workbook to produce coverage maps.

Organizations may choose to repeat Phase 2 to create multiple coverage maps (for example, to examine different implementations of a visibility surface or to detail visibility coverage provided by different products). Visibility coverage maps may take many forms, including:

- **CISA Visibility Requirements Coverage Map:** For each visibility surface, CISA may choose to create a coverage map that reflects requirements for FCEB agencies to share visibility data with CISA on an ongoing or by request basis and establish priorities for collecting and using telemetry.
- **Product Coverage Maps:** Vendors may choose to create coverage maps indicating which product tiers and configuration settings can provide visibility into the ATT&CK techniques for a given visibility surface.
- **Environment Coverage Maps:** An organization may choose to create coverage maps to understand what data is currently available to support internal cybersecurity operations or to set goals for improved visibility coverage.
- **Comparison Coverage Maps:** An FCEB Agency may create coverage maps indicating what data they plan to share with CISA to support CISA mission objectives.

In Figure 14, the visibility requirements are represented in the far left by the orange shading within the fenced in area. This represents what an organization requires to have visibility through higher fidelity observation closer to the asset or the vault and gold. The product coverage map is shown in the middle depiction by the single visibility coverage provided by the drone. Lastly, in the far right, the environment coverage map shows the visibility provided by a combination of all sensors currently deployed or planned for deployment.

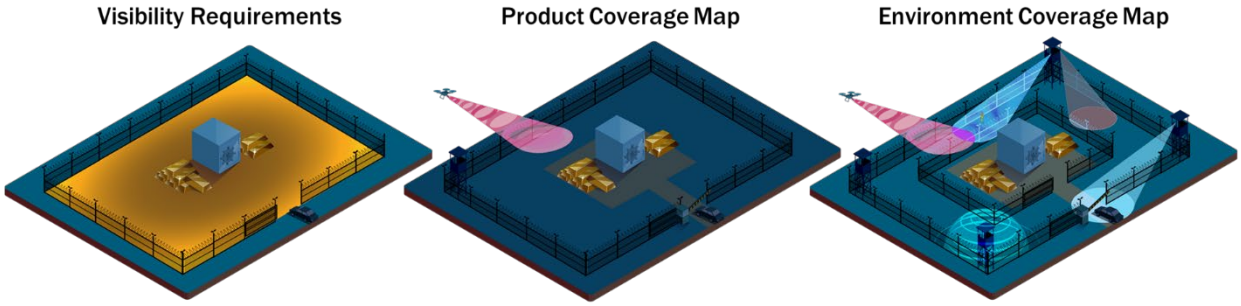

Figure 13: Visibility Requirements, Product, and Environment Coverage Maps

Figure 15 displays the environment coverage map combined with candidate product coverage maps for planning and what-if scenario consideration, thus creating a variety of coverage comparison maps. This allows organizations to evaluate the variations in visibility offered by different combinations of observation points, sensors, and products.

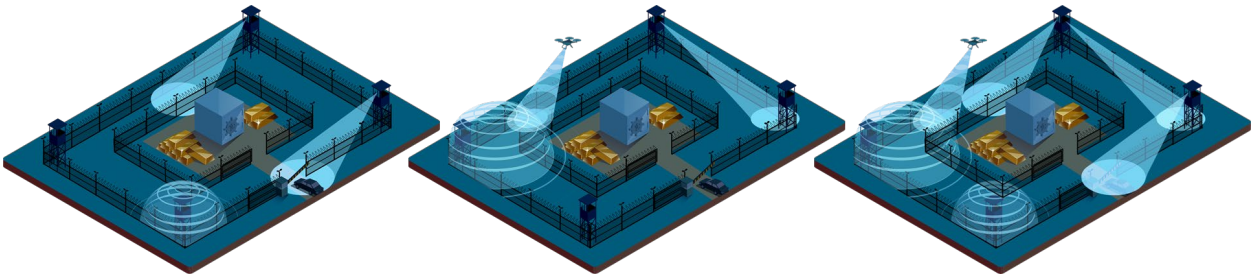

Figure 14: Coverage Comparison Map

As with earlier activities in the eVRF workflow, it may be helpful to engage a team of specialists to participate in this phase of producing visibility coverage maps. To produce accurate coverage maps and derive valuable insights, it is essential that an eVRF workbook accurately captures the technology within the environment. Include people from the organization who have expertise in configuring the relevant technologies. To create a visibility coverage map, begin with an eVRF workbook that contains a complete visibility surface definition (see Phase 1). Creation of a visibility coverage map involves four steps, shown in Figure 16:

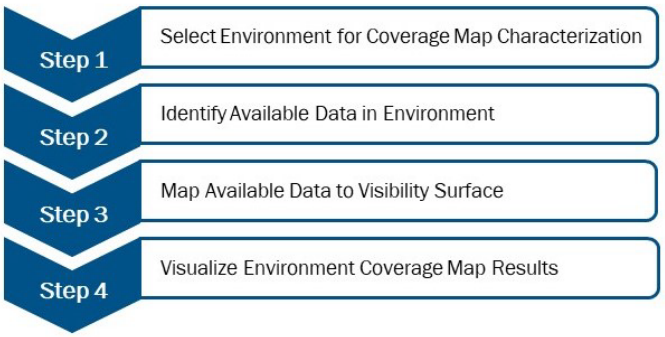

Figure 15: eVRF Workflow Phase 2

Phase 2, Step 1: Select Environment for Coverage Map Characterization

Start with the results from Phase 1, Step 5, and identify which services or applications support the visibility surface. Services identified in this step will be the basis for characterizing coverage of the entire visibility surface, so it is important to carefully consider what sources of visibility to include in the coverage maps.

Phase 2, Step 2: Identify Available Data in Environment

For each service or application, identify what logs are produced that may provide visibility into system-level and user-level events. A service or application will include an observation point with one or more sensors. For example, within the visibility surface definition for cloud business applications, relevant services and applications may include an email application, which includes mail flow logs, mailbox audit logs, etc.; an antivirus service, which includes malware protection logs; a cloud access service, which may include identity protection logs and cloud access security broker logs; and underlying cloud platform services, which include distinct event logs.

Phase 2, Step 3: Map Available Data to Visibility Surface

After identifying the available log sources, review the actual log data to verify whether the logs provide the metadata specified by the visibility surface. For each log source in the environment, enter the coverage for each piece of metadata to indicate whether the log provides that data. Continue until all log sources are addressed.

When this step is complete, the resulting work product provides detailed, application-level visibility coverage for the entire visibility surface.

Phase 2, Step 4: Visualize Environment Coverage Map Results

This step produces the coverage map. This coverage map is derived from the environment characterization provided in the previous steps of Phase 2, and it represents a summary view of the visibility coverage for all services and applications in the environment for the visibility surface. Use color-coding to indicate visibility coverage for each ATT&CK technique:

Table 1 provides the visibility coverage rubric.

Table 1: Visibility Coverage Rubric

Color	Description
N/A	Technique is not applicable to this map's scope
None	Technique is applicable but there is not visibility coverage within this map's scope
Partial	There is partial visibility coverage for the metadata events and techniques within this map's scope
Complete	There is complete visibility coverage for the metadata events and techniques within this map's scope

In the next phase of the eVRF workflow, the results provided by the Phase 2, Step 4 coverage map will be analyzed and compared with additional coverage maps to identify insights about existing coverage or answer questions related to business decisions or operational visibility.

Phase 3: Generate Visibility Coverage Comparisons for Analysis and Insights

In the final phase of the eVRF workflow, organizations create a visibility coverage comparison by combining multiple coverage maps for analysis and to generate insights. Organizations may use visibility coverage comparisons may be used to:

- Identify gaps in visibility coverage
- Establish targets for the collection of new visibility data
- Identify potential updates to system configurations
- Inform procurement decisions
- Perform "what if" scenarios prior to implementation
- Augment product offerings to provide increased breadth of visibility
- Identify redundancies or duplication of visibility

To generate valuable insights, begin with an eVRF workbook that contains both a complete visibility surface definition (see Phase 1) and a complete visibility coverage map (see Phase 2). The recommended process for generating analysis and insights from coverage map results involves the five steps displayed in Figure 17:

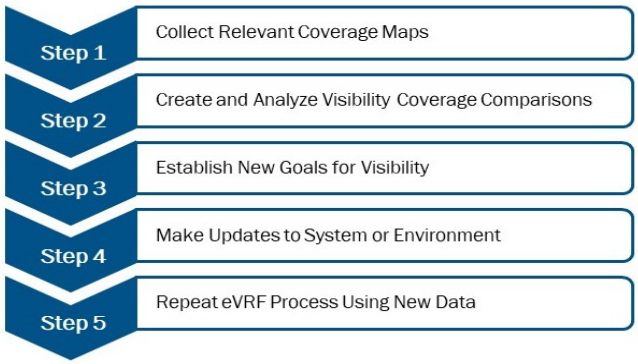

Figure 16: eVRF Workflow Phase 3

Phase 3, Step 1: Collect Relevant Coverage Maps

Collect the coverage maps to be examined or compared in the analysis. Visibility coverage comparisons allow an organization to aggregate or compare multiple coverage maps for analysis; two or more coverage maps are required for each visibility coverage comparison.

Many types of coverage maps may be available from which to choose, as described in Phase 2.

The organization performing analysis may create the coverage maps selected for this activity; or vendors or partners may provide coverage maps to support comparison or goal setting.

Organizations should customize their selection of coverage maps to suit their use case. For example, an organization seeking to understand trade-offs for an acquisition decision may choose to combine a coverage map that describes the organization's as-implemented environment with a second coverage map that describes a new product's available coverage. This would produce a visibility coverage comparison that highlights potential visibility improvements offered by the product as well as remaining gaps in coverage.

Organizations will use the coverage maps selected for this step to create a visibility coverage comparison overlay to analyze throughout the rest of the eVRF workflow.

Phase 3, Step 2: Create and Analyze Visibility Coverage Comparisons

Create one or more visibility coverage comparisons by comparing two or more coverage maps. Creating a visibility coverage comparison is currently a manual process, which may be updated and streamlined in future versions of an eVRF workbook. The easiest way to create a visibility coverage comparison currently is to arrange each coverage map side by side to compare the color-coded visibility coverage maps.

Using the visibility coverage comparison, analysts can see which ATT&CK techniques are covered by none, all, or a subset of the log sources in each individual coverage map. For organizations seeking

to compare the visibility of multiple product suites, for example, the visibility coverage comparison can show where a coverage gap exists by highlighting instances where some or none of the log sources offer visibility. Also apparent are instances where more complete visibility is offered for some products and not others.

To illustrate additional potential use cases for analysis:

- An organization may use visibility coverage comparisons to understand the effect of adding a product or service to an as-implemented environment. The comparison may illustrate redundant visibility or the need for one or more additional products to address remaining coverage gaps.
- An organization may use visibility coverage comparisons to compare an as-implemented product configuration to the optimal product configuration (e.g., as described by a vendor-provided coverage map). The comparison could inform decisions about changes to configuration settings, upgrades to products, or acquisition of new products to address coverage gaps. A department or agency may use visibility coverage comparisons to better understand the coverage provided by their as-implemented environment compared to CISA's visibility requirements. This comparison may inform decisions about new products that could address coverage gaps and mitigation strategies.
- CISA or another organization may want to use visibility coverage comparisons to compare the same visibility surface across coverage maps from many organizations. This comparison may inform decisions about new analytical toolsets or incident response activities that CISA may want to prioritize.

Phase 3, Step 3: Establish New Goals for Visibility

Organizations should set goals to improve or resolve the visibility gaps found in coverage comparisons. Some goals may also be driven by identification of redundant visibility and opportunities to improve the use of resources. In this case, organizations should consider the details of logs that appear to provide redundant coverage for the same technique—they may not be as redundant as they seem.

In addition, changes to the threat landscape, such as threat actors leveraging new techniques and sub-techniques, may initiate new visibility goals. As threat actors adopt these new approaches, organizations should work to maintain relevant visibility.

Phase 3, Step 4: Make Updates to System or Environment

In general, the solutions to visibility goals typically involve identifying new configuration settings, product upgrades, feature enhancements, or additional products or business partners that can provide the visibility desired to address gaps in coverage.

In cases of redundant or duplicative visibility or service use, the solution may be to reduce licensing, product use, or even simplify architectures.

Phase 3, Step 5: Repeat eVRF Process Using New Data

When the characterized environment is modified, the threat environment evolves, or other changes impacting the coverage maps in use occur, organizations should revise the relevant visibility surface definitions, visibility coverage maps, and visibility coverage comparisons. These eVRF components

should be living artifacts that, if re-examined regularly, can continue to provide valuable insights into an organization's current visibility posture and opportunities to improve that visibility posture.

3.2 TAILORING THE eVRF WORKFLOW

Visibility Surfaces

The workflow described in this Guidebook is intended to provide a high-level guide for organizations to adopt eVRF and to categorize visibility within their enterprise. This Guidebook defines the complete eVRF process, including the creation of a new visibility surface, which is the first step in the eVRF process. Organizations need to pursue defining and creating a visibility surface if one has not already been defined. However, in most cases, organizations will likely build on existing visibility surface definitions.

An organization can define a visibility surface if it cannot find a predefined surface that meets the organization's needs. The eVRF process includes flexibility for defining visibility surfaces and domains. To help reduce the complexity of defining a visibility surface, the organization can start by leveraging existing architectural information that has been captured regarding their enterprise (e.g., documentation created to support compliance with NIST SP 800-53[5] or other security control implementation guidance).

Prioritizing systems by considering the most important systems relative to the organization's mission or business first will provide an organization the greatest value in the near term and allows for a gradual implementation of the eVRF process throughout the organization's enterprise. Leveraging existing architectural information, along with prioritizing the most critical systems, may allow for an easier tailored approach to visibility surface generation that can be adjusted with the organization's resources in mind.

Similarly, when an organization desires to have a more comprehensive scope than an existing visibility surface, they may first inherit the pre-defined visibility surface and then expand the scope, as required.

Coverage Maps

The use of coverage maps is also integral to the eVRF process. Like visibility surfaces, organizations may be able to leverage coverage maps that outside parties have already generated. This can save an organization time and ensure alignment with other eVRF stakeholders (facilitating inter-organization analysis and comparisons). Coverage maps may be written by vendors, similar organizations, or industry leading bodies. As an organization considers whether to leverage a coverage map written by another party, they should consider multiple factors, including the map's accuracy, reliability, alignment with their selected visibility surface scope, and their specific intentions for its use.

[5] "'Security and Privacy Controls for Information Systems and Organizations," NIST, NIST Computer Security Resource Center, https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final

Rather than defining custom configurations for visibility surfaces, organizations may find it helpful to leverage established visibility surface definitions and their associated generic vendor coverage maps. They could then only consider qualities or attributes relevant to their use case. This approach would likely prove more efficient than redefining the scope of the visibility surface and then researching all candidate products and generating vendor coverage maps for each of them.

3.3 ORGANIZATION INTEGRATION OF eVRF

Identifying Visibility Gaps

The eVRF process provides a mechanism for organizations to intimately understand what visibility currently exists within their architectures. By using eVRF visibility coverage comparisons, organizations can identify visibility gaps, highlighting telemetry which might not be currently captured within the system. There might be cases where system and/or application functionality can simply be turned on or may require supplemental applications or services to achieve the desired insights. With a clear understanding of available visibility, organizations can holistically consider how they can improve their overall cybersecurity posture across their enterprise architecture.

Enhancing Resource Use Efficiency

Prioritizing security needs with limited resources is a continual challenge. The outputs that can be derived from the eVRF process and other combined efforts can aid in the prioritization of those resources. This could allow organizations to couple visibility insights with other infrastructure design elements to improve efficiency, reduce duplication of telemetry generated, and minimize redundant log collection. This deeper knowledge can also aid in driving architectural design decisions to ensure the visibility aligns with organizational goals and anticipated threats.

Supporting Zero Trust

Organizations are moving away from static network-based perimeter defenses to more dynamic and holistic security architectures. These zero trust architectures include cybersecurity protections for users, assets, data, and applications. To support the transition to zero trust architectures, CISA has developed a *Zero Trust Maturity Model.*[6] In this maturity model, "Visibility and Analytics" has been identified as a cross-cutting component for each of the cybersecurity pillars (Identity, Networks, Applications, Data, and Devices), which means that organizations are responsible for ensuring visibility coverage in each of these pillars. Organization evaluation of visibility coverage should lead to architectural improvements where coverage is lacking. This will be prevalent in an enterprise-wide scope with overlapping visibility surfaces. The organization can layer eVRF workbook coverage maps and incorporate other data into their analysis.

[6] CISA, "Zero Trust Maturity Model Pre-decisional Draft," June 2021, https://www.cisa.gov/sites/default/files/publications/CISA%20Zero%20Trust%20Maturity%20Model_Draft.pdf

4. CISA Use of eVRF

CISA's use case captured in this section describes how CISA will use eVRF with FCEB agencies and provides an example of the activities and interactions between CISA, vendors or service providers, and agencies as they work through the process described in the general workflow with this Guidebook. This section focuses on how agencies should provide visibility for government systems for CISA.

Ideally, each party will develop an iterative process between organizations to provide productive dialog and shared maturity. This will enhance the collaboration for improved feedback and refinement of requirements, products, and systems over time.

> CISA can use eVRF to define visibility requirements for FCEB agencies for select visibility surfaces.

4.1 AGENCY AND CISA BENEFITS OF eVRF

FCEB agencies will derive the benefits of eVRF mentioned in Section 1.2 as they adopt this framework. Additionally, agencies will benefit from using eVRF to meet CISA's visibility requirements in the following ways:

1. Agencies will gain better insights into their overall security posture through the enablement of enhanced visibility-informed risk analyses.
2. Agencies will be better able to analyze where gaps in visibility exist within their enterprise environment.
3. Agencies will have a greater understanding of gaps in coverage and potential risks to inform decision making processes for allocation of resources. As an agency's visibility is better understood, they will be better postured to identify and mitigate potential threats.
4. All agencies and CISA benefit from extended visibility. The inclusion of additional telemetry across domains enhances incident response and persistent hunt capabilities.
5. Using this model helps CISA aggregate and correlate threat data to aid in the timely discovery of attack campaigns facing federal enterprise systems, benefitting all agencies.
6. The frequency and availability of indicators of compromise is driven by more threat-informed and available data sets. Therefore, alignment with eVRF visibility requirements coverage maps will result in better situational awareness and availability of indicators of compromise from CISA.
7. Agencies can leverage the eVRF to help meet the logging requirements in OMB M-21-31.

4.2 ROLES AND RESPONSIBILITIES

Visibility within FCEB Agency domains helps ensure that CISA can identify threats; protect against potential attacks; and perform hunt, incident response, and analysis activities. Furthermore, this visibility enables CISA to develop and share valuable insights across the FCEB Agency domains, which provides individual agencies with valuable cybersecurity benefits. While this Guidebook does not direct or require specific actions (especially from vendors), this section generally describes the

roles and responsibilities that CISA, agencies, and vendors/service providers can take to successfully implement eVRF.

CISA Role

First and foremost, CISA is responsible for developing the eVRF and communicating guidance to other agencies, including specifications of the telemetry needs determination process and telemetry data requirements for FCEB Agency domains. CISA analyzes FCEB security events and telemetry. This responsibility guides the development of eVRF visibility requirements coverage map definitions; CISA supplies eVRF visibility surface definitions for FCEB Agency consideration on solution development and desired telemetry sharing. This ensures that the FCEB Agencies can understand the CISA eVRF visibility objectives and limitations. CISA is also responsible for updating visibility requirements to reflect changes to the threat landscape, evolution of solution offerings, FCEB Agency feedback regarding technical capabilities, and others to align to current and future telemetry needs. Figure 18 displays the CISA role in the eVRF workflow.

Figure 17: CISA Role in eVRF Workflow

Agency Role

FCEB Agencies are responsible for using eVRF-based guidance to inform their internal policies and provide alignment to agency cybersecurity needs and/or risk management planning, as appropriate. The FCEB Agencies are responsible for adopting the visibility surface definitions established by CISA and ensuring that visibility data is available to support CISA as needed. To support CISA's future potential investigative needs, agencies will need to retain and report on telemetry data on an ongoing basis. The FCEB Agencies have the responsibility to evaluate their ability to collect relevant visibility data and develop a plan to address CISA's visibility requirements. The FCEB Agencies have the responsibility to ensure configuration of telemetry generation within each domain in accordance with eVRF guidance supplied by CISA; this will ensure that the FCEB Agency security event and telemetry reported meet the CISA requirements. FCEB Agencies have a responsibility to update their as-built visibility coverage maps to reflect changes to their environment. Figure 19 displays the agency role in the eVRF workflow.

Figure 18: Agency Role in eVRF Workflow

Vendor and/or Service Provider Role

Vendors and/or service providers may elect to produce visibility coverage maps for their products and services, as well as update their visibility maps to reflect changes in their offerings over time. Figure 20 displays the vendor role in the eVRF workflow.

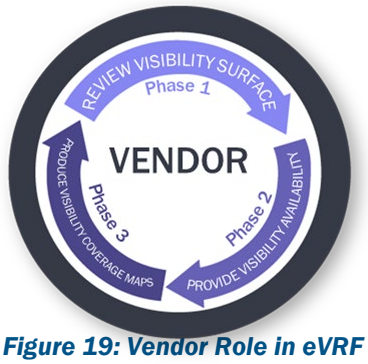

Figure 19: Vendor Role in eVRF Workflow

4.3 FCEB WORKFLOW EXAMPLE

Figure 21 shows the three entities necessary for the CISA workflow cycle. As each entity works through the process and generates data, each interacts with the other entities to refine and improve the data. Additional telemetry may become available or change over time, and all parties should update their data as technology, implementation, and target needs change.

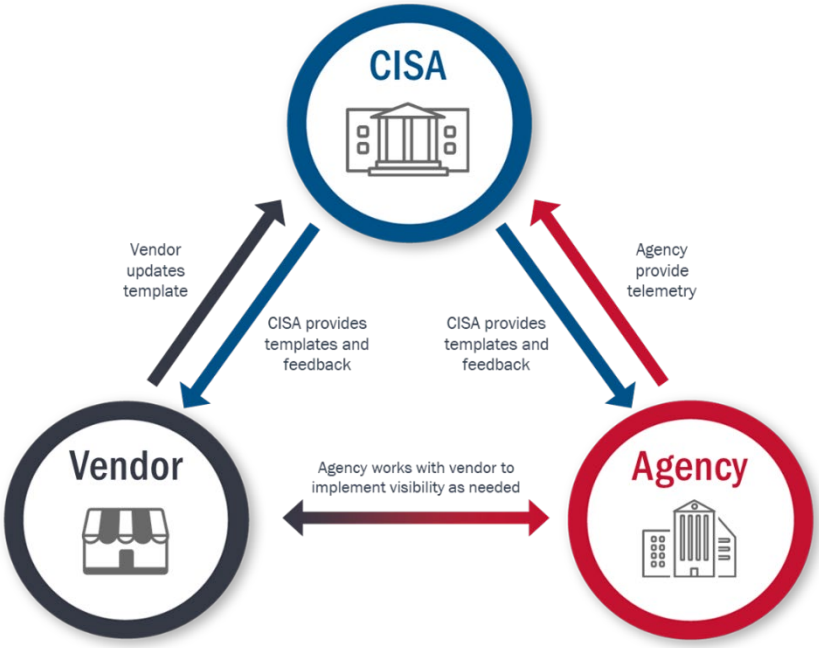

Figure 20: CISA Workflow Cycle

CISA provides its requirements for the visibility surface definition within each workbook. Agencies can work with their vendors to provide inputs for visibility telemetry that exists within each environment. Agencies provide this information to CISA for review and feedback.

To provide more information to agencies, vendors (or service providers) may want to populate the relevant data for their product offerings to identify the available telemetry. Outside of a procurement

context, vendors may elect to provide this data to CISA to allow CISA to evaluate the product's telemetry with respect to CISA requirements.

FCEB Agencies use the workbook to capture the current state of the system for existing telemetry and provide CISA with the telemetry required. Agencies may work with CISA and vendors to identify visibility gaps and determine how to improve areas with limited visibility. When one product offering might not completely provide the needed visibility, layering another product could help fill the gap. Agencies may be able to use the products and services of multiple vendors to assist with understanding potential solutions.

Workflow

The CISA workflow follows the tasks shown in Figure 22. This workflow represents tasks specific to CISA, vendors, and FCEB agencies within the previously described eVRF and specific portion alignment with the eVRF workflow.

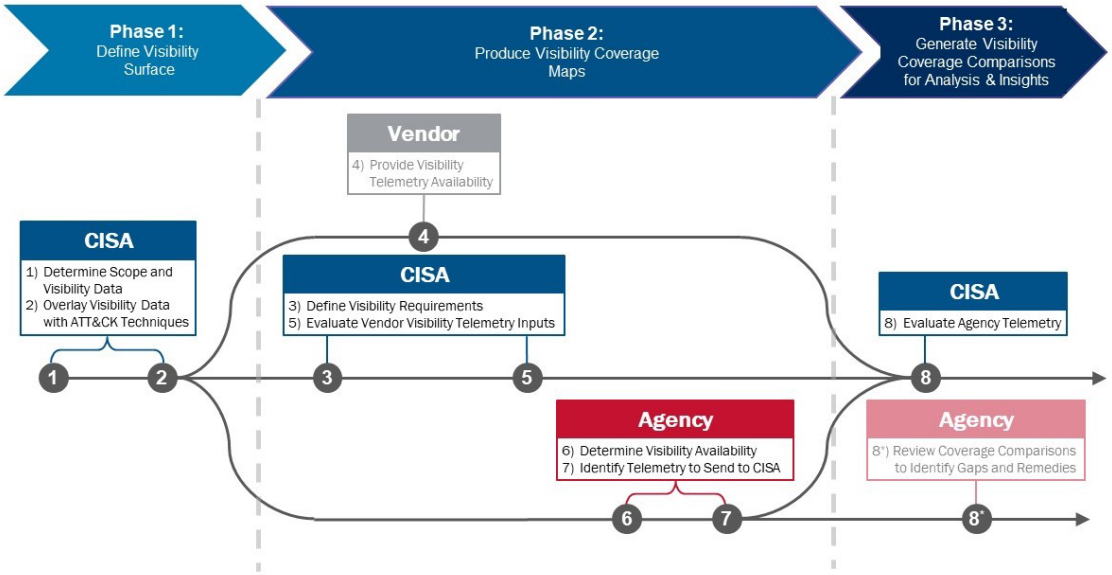

Figure 21: CISA Workflow Tasks

CISA—Phase 1, Step 1: Determine Scope of Visibility Surface

CISA will publish visibility surfaces as appropriate. For each visibility surface CISA publishes, CISA will first define the scope of the visibility surface and complete the description for the prescribed level of detail. As CISA determines which visibility surface to define next, it will consider many factors, including the log categories defined in M-21-31, Appendix C, Table 5.

CISA—Phase 1, Step 2: Identify Relevant Visibility Data

CISA determines the metadata needed for each event and category. Table 2 shows an example of a portion of the visibility surface data captured within a workbook or tool for a generic business application suite.

Table 2: Visibility Surface Example

Visibility Data			
Category	Event	Metadata	Description/Example Activities
Email	Email Received	Sender	Emails received, either internal or external to the organization
		Receiver	
		Subject	
		Other Headers	
		URLs	
		Body	
		Message Trace	
		Attachments	
	Email Sent	Sender	Emails sent, either internal or external to the organization.
		Receiver	
		Subject	
		Other Headers	
		URLs	
		Body	
		Message Trace	
		Attachments	

CISA—Phase 1, Step 3: Choose ATT&CK Matrix

With the visibility data defined, CISA then selects the desired mapping to a set of MITRE ATT&CK techniques for the given application. A MITRE ATT&CK matrix may already exist for the given product. CISA may choose to use existing matrices from MITRE, make its own, or choose a different mapping altogether with a different set of criteria. Table 3 provides a sampling of the tactics, techniques, and sub-techniques captured.

Table 3: ATT&CK Tactics, Techniques, and Sub-Techniques Example

ATT&CK Techniques		
Tactic	Technique	Sub-Technique
Initial Access	Phishing	Other Unspecified Sub-technique
		Spear phishing Link
	Valid Accounts	Other Unspecified Sub-technique
		Default Accounts
		Cloud Accounts
Persistence	Account Manipulation	Other Unspecified Sub-technique
		Exchange Email Delegate Permissions
		Add Bus. Suite Global Admin Role
	Create Account	Other Unspecified Sub-technique
		Cloud Account
	Office Application Startup	Other Unspecified Sub-technique
		Add-ins
		Office Template Macros
		Outlook Forms
		Outlook Rules
		Outlook Home Page
		Office Test
	Valid Accounts	Other Unspecified Sub-technique
		Default Accounts
		Cloud Accounts
Privilege Escalation	Valid Accounts	Other Unspecified Sub-technique
		Default Accounts
		Cloud Accounts

CISA—Phase 1, Step 4: Create ATT&CK-to-Data Overlay

CISA determines the visibility mapping to the ATT&CK techniques for the given products. For each event, CISA determines whether each event and associated data would provide visibility for each element of the ATT&CK techniques and sub-techniques within each tactic. If the metadata provides visibility, "Yes" is input into the associated field for this example. The right side of Table 4 shows the mapping.

Table 4: Overlay Visibility Data with ATT&CK Techniques Example

		ATT&CK Overlay Technique	Initial Access				
			Phishing		Valid Accounts		
		Sub-Technique	Other Unspecified Subtechnique	Spearphishing Link	Other Unspecified Subtechnique	Default Accounts	Cloud Accounts
Visibility Data							
Category	Event	Metadata					
Email	Email Received	Sender	Yes	Yes	No	No	No
		Receiver	Yes	Yes	No	No	No
		Subject	Yes	Yes	No	No	No
		Other Headers	Yes	Yes	No	No	No
		URLs	Yes	Yes	No	No	No
		Body	Yes	Yes	No	No	No
		Message Trace	Yes	Yes	No	No	No
		Attachments	Yes	Yes	No	No	No
	Email Sent	Sender	No	No	No	Yes	Yes
		Receiver	No	No	No	Yes	Yes
		Subject	No	No	No	Yes	Yes
		Other Headers	No	No	No	Yes	Yes
		URLs	No	No	No	Yes	Yes
		Body	No	No	No	Yes	Yes
		Message Trace	No	No	No	Yes	Yes
		Attachments	No	No	No	Yes	Yes

CISA—Phase 2, Step 2: Define Visibility Requirements

CISA will determine the visibility requirements for the visibility data. For instance, CISA will establish the periodicity and priority for each set of metadata for each event; and the telemetry provided to CISA will conform to these stipulated elements.

The values in the example are input as placeholders and are not intended to represent any analysis of this set of information.

Periodicity options are ongoing or by request. "Ongoing" defines metadata agencies should provide to CISA on a regular interval; CISA will negotiate details with the agencies. This telemetry will be either an automated feed or provided at a regular frequency based on the data. CISA will perform ongoing analysis on this information with advanced analytics to aid in identifying malicious activity within the agency's implemented architecture. Agencies will maintain the "by request" telemetry; as circumstances warrant, (e.g., based on analysis findings or other indicators), CISA may make a request to the agency to provide the additional information. This will allow CISA to aid the agency in performing deeper analysis in looking for additional indicators of compromise of its systems.

To help allocate resources to accommodate data requests, CISA will assign a priority level for each visibility element (the levels are 0, 1, and 2; with 0 being the highest and 2 being the lowest). CISA may consider the visibility target mapping to the visibility surface in determining the priority of each event. Based on the OMB logging requirements document, prioritization should focus on high-impact systems and high-value assets. Additional prioritization may be needed depending on an agency's specific architecture. Efforts should be focused on accommodating priority 0 requests. If a specific set of metadata cannot be provided, the agency should coordinate with CISA to implement a working solution.

The right side of Table 5 shows the visibility requirements within the workbook with the associated visibility surface information. CISA will prepopulate this portion of the workbook with CISA's requirements prior to providing the workbook to agencies. This constitutes a special purpose coverage map specific to CISAs visibility requirements regarding the subject visibility surface.

Table 5: CISA Visibility Requirements Example

Visibility Data			CISA Visibility Requirements	
Category	Event	Metadata	Ongoing or By Request	Priority
Email	Email Received	Sender	By Request	1
		Receiver	By Request	1
		Subject	By Request	1
		Other Headers	By Request	1
		URLs	By Request	1
		Body	By Request	1
		Message Trace	By Request	1
		Attachments	By Request	1
	Email Sent	Sender	By Request	1
		Receiver	By Request	1
		Subject	By Request	1
		Other Headers	By Request	1
		URLs	By Request	1
		Body	By Request	1
		Message Trace	By Request	1
		Attachments	Ongoing	0

Vendor—Phase 2, Step 1: Select Environment for Coverage Map Characterization

The vendor may identify which services or applications support the visibility surface.

Vendor—Phase 2, Step 2: Identify Available Data in Product

With the visibility mapping completed and CISA visibility requirements defined, CISA provides the workbook to the vendor (CISA—Phase 1, Step 5) to determine the visibility mapping to the ATT&CK techniques for its products. For each event, the vendor may indicate whether metadata exists that would provide visibility for each element of the ATT&CK technique and sub-technique within each tactic by writing "*Yes*" in the associated field within the workbook. As eVRF processes and tools mature, vendors may be able to provide more information, such as the level of visibility (limited/some/most).

When complete, the vendor can provide completed workbooks to CISA for adjudication. CISA will review the provided information, seek clarity on any questions, and update its processes to incorporate the vendors' input.

Table 6 provides a product visibility mapping example.

Table 6: Product Visibility Mapping Example

Product Visibility			Vendor Service	Platform Logs		Email Application			
			Telemetry Source	Event Log (Tier 1)	Event Log (Tier 2)	Mailbox Logs (Tier 1)	Mailbox Logs (Tier 2)	Mail Flow Logs (Tier 2)	Phishing Protections (Tier 2)
Visibility Data									
Category	Event	Metadata							
Email	Email Received	Sender				Yes	Yes		
		Receiver				Yes	Yes		
		Subject				Yes	Yes		
		Other Headers				Yes	Yes		
		URLs					Yes		
		Body					Yes		
		Message Trace						Yes	
		Attachments					Yes		
	Email Sent	Sender							
		Receiver							
		Subject							
		Other Headers							
		URLs							
		Body							
		Message Trace							
		Attachments							Yes

CISA—Phase 1, Step 5: Create Data Entry Template for Coverage Maps

CISA may use vendor submitted information, if any, to update and improve CISA's visibility requirements. CISA then creates the templates for Phase 2, which agencies will use to characterize their environment and identify the visibility data that is available.

Agency—Phase 2, Step 1: Select Environment for Coverage Map Characterization

The agency identifies which services or applications support the visibility surface.

Agency—Phase 2, Steps 2 & 3: Identify & Map Available Data to Visibility Surface

The agency will use the workbook to determine the visibility available within its system architecture. This will be a review of all the agency's systems and vendor products associated with each visibility surface. Within each domain, the agency will determine what observation points and sensors are deployed, what telemetry they have available, and how data is currently being captured. If the agency is not currently collecting the telemetry, the agency should consider efforts to refine the architecture to either capture the telemetry or provide a CISA-approved alternative. In cases where the agency is unable to provide telemetry, agency personnel should work with CISA on an agreed path forward.

Identify Telemetry to Send to CISA

Based upon the agency's architecture and visibility determination, the agency will identify appropriate telemetry associated with visibility surfaces and requirements to send to CISA.

FCEB Agencies may want to refer to CISA Visibility Requirements Coverage Maps for agencies to provide visibility data. For each piece of visibility data, the table includes a priority ranking and

indicates whether the data should be provided on an ongoing basis or whether the data should be available to CISA upon request. Agencies may also refer to *NCPS Cloud Interface Reference Architecture Volumes 1 and 2* for details regarding telemetry generation, processing, and reporting to CISA.

Agency—Phase 2, Step 4: Visualize Environment Coverage Map Results

Agencies can generate coverage maps specific to the agency as part of Phase 2. The telemetry availability as implemented is based on the product software selection and mapped to the ATT&CK overlay. This will represent any gaps that exist due to applications that either are not used or have not been fully implemented.

In the example map shown in Figure 23, techniques where implemented coverage is not present when CISA has specified a tie to the event, metadata for the technique will be shaded yellow to indicate the deviation from the defined visibility surface.

ATT&CK Matrix Coverage Map for Characterized Environment Compared to Visibility Surface

Initial Access	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Collection	Impact
Phishing 73%	Account Manipulation 75%	Valid Accounts 47%	Impair Defenses 79%	Brute Force 60%	Account Discovery 55%	Internal Spearphishing 54%	Data from Information Repositories 69%	Endpoint Denial of Service 100%
Valid Accounts 50%	Create Account 75%		Use Alternate Authentication Material 47%	Forge Web Credentials 73%	Cloud Service Dashboard 67%	Use Alternate Authentication Material 41%	Email Collection 0%	Network Denial of Service 100%
	Office Application Startup 71%		Valid Accounts 47%	Steal Application Access Token 78%	Cloud Service Discovery 50%			
	Valid Accounts 47%			Steal Web Session Cookie	Permission Groups Discovery 50%			
				Unsecured Credentials 67%	Software Discovery 60%			

Figure 22: Coverage Comparison—Environment to Product [7]

CISA—Phase 3: Analyze Agency Telemetry

CISA will check telemetry it receives to ensure alignment with information the agency specifies. CISA will work with the agency to resolve any issues in transmission. This will be an ongoing verification to ensure that the telemetry aligns to the specifications.

CISA will then review the telemetry inputs the agency provides to identify discrepancies and work with the agency to resolve gaps in telemetry to ensure CISA receives the required visibility telemetry. This will be a recurring process as the agency updates the information available to provide to CISA and makes modifications to its architecture.

[7] The percentages in the figure represent the fraction of the required telemetry within the visibility surface that is satisfied by the organization's identified telemetry availability for each technique.

Agency—Phase 3: Generate Visibility Coverage Comparisons for Analysis and Insights

The coverage maps generated in Phase 2 will aid in building broader visibility coverage comparisons to identify opportunities to close those gaps with other products.

4.4 FCEB USE OF VISIBILITY COVERAGE COMPARISONS

Each stakeholder will use visibility coverage comparisons differently. Fundamentally, each visibility coverage comparison shows a summation of coverage maps, representing the telemetry of multiple products or services.

The versatility of visibility coverage comparisons allows for the evaluation of both agencies' telemetry as well as telemetry sources as offered by vendors. Visibility coverage comparisons can be used to compare the relative strength of two different collections of services or two different agencies. Ultimately, wide usage of visibility coverage comparisons will provide more informed decision-making to maximize breadth and depth of telemetry coverage across the FCEB.

Agencies

There are multiple ways an agency can use visibility coverage comparisons to analyze and improve its defensive posture. An agency should internally maintain a comprehensive version of a single or multiple visibility coverage comparisons to evaluate its cybersecurity posture as an organization or in a division. A simple analysis of the agency's visibility coverage comparison quickly conveys gaps and overlaps in telemetry. Each agency will provide CISA with a visibility coverage comparison with at least the minimum required details on ATT&CK framework coverage. An agency can, by itself or in coordination with CISA, compare its current visibility coverage comparison with CISA's recommended telemetry coverage.

Gaps in an agency's visibility coverage comparisons reflect gaps in the implementing agency's applications. Some may be covered under CISA's supported service offerings. CISA recommends using its eVRF insights into visibility to determine which products strengthen the agency's overall posture the most. Even if CISA's federal service offerings overlap with an agency's current coverage, using CISA's services may still provide benefits such as federated threat sharing or more seamless integration with other CISA-supported offerings.

Agencies will not export all available data to CISA, but working through this process will help an agency identify where it has telemetry available to provide insights into potential malicious activity on its networks. This process will also aid in identifying where gaps may exist through the generated coverage maps within the workbooks, as well as identify potential areas where an agency may want to focus on shoring up their architectures.

The visibility coverage comparison(s) represent the difference between what is available within the product suite and what has been implemented. Techniques where implemented coverage is not present when the product suite does provide telemetry options will be shaded red to indicate the agency may have a mechanism to fill those gaps.

CISA

CISA will use visibility coverage comparisons to empirically inform the list of telemetry it requires agencies to provide to CISA on an ongoing or on-demand basis. Organizations within CISA will use visibility coverage comparisons to inform which telemetry to prioritize in analytical toolsets or incident response engagement activities. Over time, CISA will build recommended visibility coverage comparisons based on profiles of certain combinations of products or services. CISA will be able to show differences between ideal visibility coverage comparisons and an agency's current visibility coverage comparison, as well as make specific mitigation recommendations.

Vendors

Vendors of enterprise business applications or cloud security software will benefit from clear security evaluation criteria of their products. Leveraging the methodology outside of a procurement context, vendors can complete eVRF workbooks to describe the telemetry their products and services make available.

As part of the iterative and ongoing improvement process, vendors can work with CISA to determine how to satisfy information needs in cases where metadata are unavailable.

5. Conclusion

The eVRF defines the concepts, requirements, and mechanisms for CISA, FCEB Agencies, and other partners to identify, characterize, collect, and apply visibility data to mitigate threats. The eVRF uses multiple work products to define and describe key concepts, roles and responsibilities, and workflows; identifies mechanisms to define a visibility surface; and enables organizations to produce their own visibility coverage maps and visibility coverage comparisons.

Appendix A: Relationship of eVRF to CISA Programs

This appendix describes the relationship between eVRF and other CISA programs.

TRUSTED INTERNET CONNECTIONS (TIC)

The goal of the Trusted Internet Connections (TIC) | CISA (TIC) initiative[8] is to secure federal data, networks, and boundaries; and to provide visibility into agency "traffic," including both network traffic and traffic between designated trust zones in a particular use case. The scope of TIC includes cloud, mobile, encrypted applications, services, and environments; therefore, this overlaps with the scope of eVRF. TIC use cases provide guidance on the implementation of security capabilities, but this guidance does not prescribe what telemetry an agency should collect and maintain. Additionally, a TIC use case identifies *where* CISA telemetry may be required for the use case; however, a use case does not identify *what* telemetry is required.

NATIONAL CYBERSECURITY PROTECTION SYSTEM (NCPS)

The National Cybersecurity Protection System (NCPS)[9] is an integrated system-of-systems that delivers a range of capabilities such as intrusion detection, analytics, information sharing, and intrusion prevention. These capabilities provide a technological foundation that enables CISA to secure and defend the FCEB agencies' information technology infrastructure against advanced cyber threats.

The NCPS Program is evolving to ensure that security information about cloud-based traffic can be captured and analyzed. To support this goal, CISA is piloting a cloud-based architecture, the Cloud Log Aggregation Warehouse (CLAW), to collect and analyze agency cloud security data. The *NCPS Cloud Interface Reference Architecture* (NCIRA) explains how agencies can create reporting patterns to describe their process for providing cloud-generated security information to CLAW. The reporting pattern has an attribute for "telemetry type," with several options to categorize common types of cloud telemetry. The NCIRA documents describe multiple options for sharing cloud telemetry with CISA but do *not* define specific requirements for what cloud telemetry is shared. CISA will use eVRF as a framework to define telemetry requirements.

[8] https://www.cisa.gov/trusted-internet-connections

[9] "National Cybersecurity Protection System," cisa.gov, CISA, https://www.cisa.gov/resources-tools/programs/national-cybersecurity-protection-system.

CONTINUOUS DIAGNOSTICS AND MITIGATION (CDM)

The Continuous Diagnostics and Mitigation (CDM)[10] program provides cybersecurity tools, integration services, and dashboards to participating agencies to support them in improving their respective security posture. Future CDM requirements may specify collection of internal telemetry in accordance with Section 7(f) of Executive Order 14028.[11] There may be overlap of CDM telemetry with the CISA Visibility Requirements.

GOVCAR (CYBERSECURITY ARCHITECTURE REVIEW)

CISA uses the .govCAR methodology to conduct threat-based assessments of cyber capabilities for the Federal Civilian Executive Branch (.gov domain). Viewing a target architecture the way a threat actor would, provides a threat-informed approach to identify where mitigations could be applied to provide the best defense against all phases of a cyberattack. Similarly, eVRF is a threat-based framework for identifying visibility data that can address adversarial attacks.

[10] "Continuous Diagnostics and Mitigation (CDM) Program," cisa.gov, CISA, https://www.cisa.gov/resources-tools/programs/continuous-diagnostics-and-mitigation-cdm-program.

[11] Executive Order 14026, "*Improving the Nation's Cybersecurity*", (May 2021). https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/.

Appendix B: Key Terms

The eVRF uses key terms, which are summarized here for reference.

Term	Description
Category	A component (application, software, service, etc.) of a system in which cyber-observable data exists.
Coverage Map	See Visibility Coverage Map(s).
Cyber Observable Data	The data elements or artifacts, e.g., configurations or configuration settings, data flows, logs, packet data, etc., which describe an event (benign or malicious) or the state on a network or system, and which contribute to a visibility surface.
Domain	A platform specific environment, e.g., cloud, mobile, on-site, which may represent a component of the cybersecurity scope within an agency's modern enterprise.
Event	A process that occurs within a defined component (of a visibility surface).
Metadata	The data or information elements (within a visibility surface) that documents the state of a system, that an event occurred, and/or how it may have occurred.
Observation Point	An observation point defines the architecture location of a telemetry source in the given domain.
Telemetry	Artifacts derived from security capabilities that provide visibility into security posture, often through automated collections.
TTPs	Threat actor tactics, techniques, and procedures (TTPs); typically, as they relate to visibility surfaces that may enable an organization to identify them.
Visibility	Visibility provides context-specific insights about the activity taking place within a given environment. CISA uses the term visibility to refer to the observable artifacts of digital events, and the characteristics of the digital environment in which those events take place.
Visibility Coverage Map	A visibility coverage map characterizes the ability of a product or organization to sufficiently address a visibility surface through available cyber-observable data.
Visibility Coverage Comparison	Overlay of one or more coverage maps applied simultaneously to the MITRE ATT&CK framework to better understand the holistic cybersecurity posture of a group of deployed products and services.
Visibility Surface	A visibility surface refers to a digital environment for which cyber-observable data exists or should exist. A visibility surface is made up of many different points from which a system can be observed and describes the scope of the environment and its relevant data and TTPs.

Appendix C: Key Documents

The eVRF leverages concepts presented in several key documents sets.

MITRE ATT&CK

The MITRE ATT&CK framework[12] categorizes the tactics, and techniques, and procedures (TTPs) threat actors employ in compromising computing infrastructure. Tactics refer to objectives a threat actor tries to achieve, and techniques refer to how a threat actor pursues a given tactic. ATT&CK has been specialized for cloud environments in the form of the Cloud Matrix. The matrix identifies 11 tactics: initial access, persistence, privilege escalation, defense evasion, credential access, discovery, lateral movement, collection, exfiltration, execution and impact. Each of these tactics consists of individual techniques that a threat actor may employ and that often result in visible signs contained in telemetry. eVRF will provide mappings between telemetry and the visibility they provide on threat actor tactics and techniques.

NCPS CLOUD INTERFACE REFERENCE ARCHITECTURE

The *NCPS Cloud Interface Reference Architecture*[13] [14] ("NCIRA") provides a framework of "reporting patterns" that agencies can use to send cloud telemetry to CISA. Each reporting pattern consists of choices around how telemetry is generated, how telemetry is processed, and how telemetry is delivered to CISA. NCIRA is therefore a guide for agencies on *how* to share telemetry with CISA and eVRF is a guide for *what* telemetry agencies should share.

ZERO TRUST MATURITY MODEL

CISA has released a *Zero Trust Maturity Model*[15] in response to the Executive Order 14028, *Improving the Nation's Cybersecurity.*[16] The maturity model describes a gradient of implementation across five distinct pillars: Identity, Device, Network, Application Workload, and Data. The maturity model includes very high-level guidance regarding "Visibility and Analytics" for each pillar. Agencies

---

[12] "MITRE ATT&CK," MITRE ATT&CK, 2015–2022, https://attack.mitre.org/.

[13] CISA, "National Cybersecurity Protection System (NCPS) Cloud Interface Reference Architecture, Volume 1 - General Guidance" July 24, 2020, https://www.cisa.gov/sites/default/files/publications/CISA_NCPS_Cloud_Interface_RA_Volume-1.pdf.

[14] CISA, "National Cybersecurity Protection System Cloud Interface Reference Architecture, Volume 2: Reporting Pattern Catalogue,"May 14, 2021, https://www.cisa.gov/sites/default/files/publications/NCPS%20Cloud%20Interface%20RA%20Volume%20Two%202021-06-11%20%28508%20COMPLIANT%29.pdf.

[15] CISA, "Zero Trust Maturity Model Pre-decisional Draft," June 2021, https://www.cisa.gov/sites/default/files/publications/CISA%20Zero%20Trust%20Maturity%20Model_Draft.pdf.

[16] Executive Order 14026, "*Improving the Nation's Cybersecurity*", (May 2021). https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/.

can use eVRF to continually incorporate visibility as they evolve their zero trust architectures over time.

OMB MEMO M-21-31

OMB has released a memorandum[17] ("Improving the Federal Government's Investigative and Remediation Capabilities Related to Cybersecurity Incidents") on logging, log retention, and log management for FCEB Agencies in support of the *Executive Order on Improving the Nation's Cybersecurity*.[18] The memo includes a maturity model for event log management and logging requirements for many log categories across an enterprise.

CLOUD TECHNICAL REFERENCE ARCHITECTURE

The *Cloud Security Technical Reference Architecture (TRA)*[19] provides strategic and technical guidance to agencies as they adopt cloud technology. The TRA focused on shared services, designing software in the cloud, and cloud security posture management (CSPM). The CSPM discussion includes considerations for visibility and sensor positioning, and cloud telemetry and logs.

OMB MEMO M-22-09

OMB Memo M-22-09, the *OMB Zero Trust Strategy*[20] ("Moving the U.S. Government Toward Zero Trust Cybersecurity Principles"), clarifies priorities for federal civilian agencies as they transition to zero trust architectures. The strategy recognizes that this is a paradigm shift for agencies and that agencies and CISA must have visibility beyond an agency's perimeter. Enterprise-wide logging is a key component to how agencies deploy zero trust architectures.

---

[17] Acting Director Shalanda D. Young to the Heads of Executive Departments and Agencies, August 27, 2021, Executive Office of the President Office of Management and Budget, https://www.whitehouse.gov/wp-content/uploads/2021/08/M-21-31-Improving-the-Federal-Governments-Investigative-and-Remediation-Capabilities-Related-to-Cybersecurity-Incidents.pdf, M-21-31,.

[18] Executive Order 14026, "*Improving the Nation's Cybersecurity*", (May 2021). https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/.

[19] "Cloud Security Technical Reference Architecture," cisa.gov, CISA, October 1, 2021, https://www.cisa.gov/publication/cloud-security-technical-reference-architecture.

[20] "Federal Zero Trust Strategy," zerotrust.cyber.gov, OMB and CISA, https://zerotrust.cyber.gov/federal-zero-trust-strategy/.