NIST SPECIAL PUBLICATION 1800-11B

Data Integrity

Recovering from Ransomware and Other Destructive Events

Volume B:
Approach, Architecture, and Security Characteristics

Timothy McBride
National Cybersecurity Center of Excellence
National Institute of Standards and Technology

Michael Ekstrom
Lauren Lusty
Julian Sexton
Anne Townsend
The MITRE Corporation
McLean, VA

September 2017

DRAFT

This publication is available free of charge from:
https://nccoe.nist.gov/projects/building-blocks/data-integrity

DISCLAIMER

Certain commercial entities, equipment, products, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by NIST or NCCoE, nor is it intended to imply that the entities, equipment, products, or materials are necessarily the best available for the purpose.

National Institute of Standards and Technology Special Publication 1800-11b, Natl. Inst. Stand. Technol. Spec. Publ. 1800-11b, 64 pages, (September 2017), CODEN: NSPUE2

FEEDBACK

You can improve this guide by contributing feedback. As you review and adopt this solution for your own organization, we ask you and your colleagues to share your experience and advice with us.

Comments on this publication may be submitted to di-nccoe@nist.gov.

Public comment period: September 6, 2017 through November 6, 2017

All comments are subject to release under the Freedom of Information Act (FOIA).

National Cybersecurity Center of Excellence
National Institute of Standards and Technology
100 Bureau Drive
Mailstop 2002
Gaithersburg, MD 20899
Email: nccoe@nist.gov

NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

The National Cybersecurity Center of Excellence (NCCoE), a part of the National Institute of Standards and Technology (NIST), is a collaborative hub where industry organizations, government agencies, and academic institutions work together to address businesses' most pressing cybersecurity issues. This public-private partnership enables the creation of practical cybersecurity solutions for specific industries, as well as for broad, cross-sector technology challenges. Through consortia under Cooperative Research and Development Agreements (CRADAs), including technology partners—from Fortune 50 market leaders to smaller companies specializing in IT security—the NCCoE applies standards and best practices to develop modular, easily adaptable example cybersecurity solutions using commercially available technology. The NCCoE documents these example solutions in the NIST Special Publication 1800 series, which maps capabilities to the NIST Cyber Security Framework and details the steps needed for another entity to recreate the example solution. The NCCoE was established in 2012 by NIST in partnership with the State of Maryland and Montgomery County, Md.

To learn more about the NCCoE, visit https://nccoe.nist.gov. To learn more about NIST, visit https://www.nist.gov.

NIST CYBERSECURITY PRACTICE GUIDES

NIST Cybersecurity Practice Guides (Special Publication Series 1800) target specific cybersecurity challenges in the public and private sectors. They are practical, user-friendly guides that facilitate the adoption of standards-based approaches to cybersecurity. They show members of the information security community how to implement example solutions that help them align more easily with relevant standards and best practices and provide users with the materials lists, configuration files, and other information they need to implement a similar approach.

The documents in this series describe example implementations of cybersecurity practices that businesses and other organizations may voluntarily adopt. These documents do not describe regulations or mandatory practices, nor do they carry statutory authority.

ABSTRACT

Businesses face a near-constant threat of destructive malware, ransomware, malicious insider activities, and even honest mistakes that can alter or destroy critical data. These data corruption events could cause a significant loss to a company's reputation, business operations, and bottom line.

These types of adverse events, that ultimately impact data integrity, can compromise critical corporate information including emails, employee records, financial records, and customer data. It is imperative for organizations to recover quickly from a data integrity attack and trust the accuracy and precision of the recovered data.

The National Cybersecurity Center of Excellence (NCCoE) at NIST built a laboratory environment to explore methods to effectively recover from a data corruption event in various Information Technology (IT) enterprise environments. NCCoE also implemented auditing and reporting IT system use to support incident recovery and investigations.

This NIST Cybersecurity Practice Guide demonstrates how organizations can implement technologies to take immediate action following a data corruption event. The example solution outlined in this guide encourages effective monitoring and detection of data corruption in standard, enterprise components as well as custom applications and data composed of open-source and commercially available components.

KEYWORDS

business continuity; data integrity; data recovery; malware; ransomware

ACKNOWLEDGMENTS

We are grateful to the following individuals for their generous contributions of expertise and time.

Name	Organization
Steve Petruzzo	GreenTec USA
Steve Roberts	Hewlett Packard Enterprise
Dave Larimer	IBM Corporation
John Unthank	IBM Corporation
Jim Wachhaus	Tripwire
Donna Koschalk	Veeam Software Corporation
Brian Abe	The MITRE Corporation
Sarah Kinling	The MITRE Corporation
Josh Klosterman	The MITRE Corporation
Susan Urban	The MITRE Corporation

Name	Organization
Mary Yang	The MITRE Corporation

The Technology Partners/Collaborators who participated in this build submitted their capabilities in response to a notice in the Federal Register. Respondents with relevant capabilities or product components were invited to sign a Cooperative Research and Development Agreement (CRADA) with NIST, allowing them to participate in a consortium to build this example solution. We worked with:

Technology Partner/Collaborator	Build Involvement
GreenTec USA	GreenTec WORMdisk, v151228
Hewlett Packard Enterprise	HPE ArcSight ESM, v6.9.1 HPE ArcSight Connector, v7.4.0
IBM Corporation	IBM Spectrum Protect, v8.1.0
Tripwire	Tripwire Enterprise, v8.5 Tripwire Log Center, v7.2.4.80
Veeam Software Corporation	Veeam Availability Suite, v9.5

Contents

List of Figures

List of Tables

1 Summary

Businesses face a near-constant threat of destructive malware, ransomware, malicious insider activities, and even honest mistakes that can alter or destroy critical data. These types of adverse events ultimately impact data integrity (DI). It is imperative for organizations to recover quickly from a DI attack and trust the accuracy and precision of the recovered data.

The National Cybersecurity Center of Excellence (NCCoE) at the National Institute of Standards and Technology (NIST) built a laboratory environment to explore methods to recover from a data corruption event in various information technology (IT) enterprise environments. The example solution outlined in this guide describes the solution built in the NCCoE lab. It encourages effective monitoring and detection of data corruption in standard enterprise components as well as custom applications and data composed of open-source and commercially available components.

The goals of this NIST Cybersecurity Practice Guide are to help organizations confidently:

- restore data to its last known good configuration
- identify the correct backup version (free of malicious code and data for data restoration)
- identify altered data as well as the date and time of alteration
- determine the identity/identities of those who alter data
- identify other events that coincide with data alteration
- determine any impact of the data alteration

For ease of use, here is a short description of the different sections of this volume.

- **Section 1:** *Summary* presents the challenge addressed by the NCCoE project, with an in-depth look at our approach, the architecture, and the security characteristics we used; the solution demonstrated to address the challenge; benefits of the solution; and the technology partners that participated in building, demonstrating, and documenting the solution. The Summary also explains how to provide feedback on this guide.
- **Section 2:** *How to Use This Guide* explains how readers—business decision makers, program managers, and IT professionals (e.g., systems administrators)—might use each volume of the guide.
- **Section 3:** *Approach* offers a detailed treatment of the scope of the project and describes the assumptions on which the security platform development was based, the risk assessment that informed platform development, and the technologies and components that industry collaborators gave us to enable platform development.

- **Section 4:** *Architecture* describes the usage scenarios supported by project security platforms, including Cybersecurity Framework [1] functions supported by each component contributed by our collaborators.
- **Section 5:** *Example Implementation* provides an in-depth description of the implementation developed in the NCCoE's lab environment.
- **Section 6:** *Security Characteristics Analysis* provides details about the tools and techniques we used to perform risk assessments.
- **Section 7:** *Functional Evaluation* summarizes the test sequences we employed to demonstrate security platform services, the Cybersecurity Framework functions to which each test sequence is relevant, and the NIST Special Publication (SP) 800-53-4 controls that applied to the functions being demonstrated.
- **Section 8:** *Future Build Considerations* is a brief treatment of other DI implementations NIST is considering consistent with Framework Core Functions: Identify, Protect, Detect and Respond, System Level Recovery, and Dashboarding.

1.1 Challenge

Thorough collection of quantitative and qualitative data is important to organizations of all types and sizes. It can impact all aspects of a business, including decision making, transactions, research, performance, and profitability. When these data collections sustain a DI attack caused by unauthorized insertion, deletion, or modification of information, it can impact emails, employee records, financial records, and customer data, rendering it unusable or unreliable. Some organizations have experienced systemic attacks that caused a temporary cessation of operations. One variant of a DI attack—ransomware—encrypts data and holds it hostage while the attacker demands payment for the decryption keys.

When DI events occur, organizations must be able to recover quickly from the events and trust that the recovered data is accurate, complete, and free of malware.

1.2 Solutions

The NCCoE implemented a solution that incorporates appropriate actions in response to a detected DI event. The solution is comprised of multiple systems working together to recover from a data corruption event in standard enterprise components. These components include, but are not limited to, mail servers, databases, end user machines, virtual infrastructure, and file share servers. Essential to the recovery is an investigation into auditing and reporting records to understand the depth and breadth of the event across these systems and inclusive of user activity.

The NCCoE sought existing technologies that provided the following capabilities:

- secure storage
- logging
- virtual infrastructure
- corruption testing
- backup capability

While the NCCoE used a suite of commercial products to address this cybersecurity challenge, this guide does not endorse any particular products—nor does it guarantee compliance with any regulatory initiatives. Your organization's information security experts should identify the products that will best integrate with your existing tools and IT system infrastructure. Your organization can adopt this solution or one that adheres to these guidelines in whole, or you can use this guide as a starting point for tailoring and implementing parts of the solution. In developing our solution, we used standards and guidance from the following, which can also provide your organization relevant standards and best practices:

- NIST Framework for Improving Critical Infrastructure Cybersecurity (commonly known as the NIST CSF) [1]
- NISTIR 8050: Executive Technical Workshop on Improving Cybersecurity and Consumer Privacy [2]
- Special Publication 800-30 Rev. 1: Guide for Conducting Risk Assessments [3]
- Special Publication 800-37 Rev. 1: Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Lifecycle Approach [4]
- Special Publication 800-39: Managing Information Security Risk [5]
- Special Publication 800-40 Rev. 3: Guide to Enterprise Patch Management Technologies [6]
- Special Publication 800-53 Rev. 4: Security and Privacy Controls for Federal Information Systems and Organizations [7]
- FIPS 140-2: Security Requirements for Cryptographic Modules [8]
- Special Publication 800-86: Guide to Integrating Forensic Techniques into Incident Response [9]
- Special Publication 800-92: Guide to Computer Security Log Management [10]
- Special Publication 800-100: Information Security Handbook: A Guide for Managers [11]
- Special Publication 800-34 Rev. 1: Contingency Planning Guide for Federal Information Systems [12]
- Office of Management and Budget, Circular Number A-130: Managing Information as a Strategic Resource [13]

- Special Publication 800-61 Rev. 2: Computer Security Incident Handling Guide [14]
- Special Publication 800-83 Rev. 1: Guide to Malware Incident Prevention and Handling for Desktops and Laptops [15]
- Special Publication 800-150: Guide to Cyber Threat Information Sharing [16]
- Special Publication 800-184: Guide for Cybersecurity Event Recovery [17]

1.3 Benefits

The NCCoE's practice guide can help your organization:

- develop an implementation plan for recovering from a cybersecurity event
- facilitate a smoother recovery from an adverse event and maintain operations
- maintain integrity and availability of data that is critical to supporting business operations and revenue-generating activities
- manage enterprise risk (consistent with the foundations of the NIST CSF)

2 How to Use This Guide

This NIST Cybersecurity Practice Guide demonstrates a standards-based reference design and provides users with the information they need to replicate a solution to recover from attacks on DI to a last known good. This reference design is modular and can be deployed in whole or in part.

This guide contains three volumes:

- NIST SP 1800-11a: *Executive Summary*
- NIST SP 1800-11b: *Approach, Architecture, and Security Characteristics* – what we built and why **(you are here)**
- NIST SP 1800-11c: *How-To Guides* – instructions for building the example solution

Depending on your role in your organization, you might use this guide in different ways.

Business decision makers, including chief security and technology officers, will be interested in the *Executive Summary (NIST SP 1800-11a)*, which describes the:

- challenges enterprises face in attacks on DI
- example solution built at the NCCoE
- benefits of adopting the example solution

Technology or security program managers who are concerned with how to identify, understand, assess, and mitigate risk will be interested in this part of the guide, *NIST SP 1800-11b,* which describes what we did and why. The following sections will be of particular interest:

- Section 3.4.1, Assessing Risk Posture - describes the risk analysis we performed.
- Section 3.4.2, Security Control Map - maps the security characteristics of this example solution to cybersecurity standards and best practices.

You might share the *Executive Summary, NIST SP 1800-11a,* with your leadership team members to help them understand the importance of adopting standards-based methods to recover from attacks on DI to a last known good.

IT professionals who want to implement a similar approach will find the whole practice guide useful. You can use the "how-to" portion of the guide, *NIST SP 1800-11c*, to replicate all or parts of the build created in our lab. The guide provides specific product installation, configuration, and integration instructions. We do not recreate the product manufacturers' documentation, which is generally widely available. Rather, we show how we incorporated the products together in our environment to create an example solution.

This guide assumes that IT professionals have experience implementing security products within the enterprise. While we used a suite of commercial products, this guide does not endorse these particular products. Your organization can adopt this solution or one that adheres to these guidelines in whole, or you can use this guide as a starting point for tailoring parts of it to recover from attacks on DI. Your organization's security experts should identify the products that will best integrate with your existing tools and IT system infrastructure. We hope you will seek products that are congruent with applicable standards and best practices. Section 3.5, Technologies, lists the products we used and maps them to the cybersecurity controls provided by this reference solution.

A NIST Cybersecurity Practice Guide does not describe "the" solution, but a possible solution. This is a draft guide. We seek feedback on its contents and welcome your input. Comments, suggestions, and success stories will improve subsequent versions of this guide. Please contribute your thoughts to di-nccoe@nist.gov.

2.1 Typographic Conventions

The following table presents typographic conventions used in this volume.

Typeface/ Symbol	Meaning	Example
Italics	filenames and pathnames references to documents that are not hyperlinks, new terms, and placeholders	For detailed definitions of terms, see the *NCCoE Glossary.*
Bold	names of menus, options, command buttons and fields	Choose **File > Edit**.
`Monospace`	command-line input, on-screen computer output, sample code examples, status codes	`mkdir`
`Monospace Bold`	command-line user input contrasted with computer output	**`service sshd start`**
blue text	link to other parts of the document, a web URL, or an email address	All publications from NIST's National Cybersecurity Center of Excellence are available at http://nccoe.nist.gov

3 Approach

Based on key points expressed in *NIST IR 8050: Executive Technical Workshop on Improving Cybersecurity and Consumer Privacy* (2015) [2], the NCCoE is pursuing a series of DI projects to map the core functions of the NIST Cybersecurity Framework. This initial project is centered on the core function of recovery, which is focused on recovering data to the last known good state. NCCoE engineers working with a Community of Interest (COI) defined the requirements for the DI project.

Members of the COI, which include participating vendors referenced in this document, contributed to the development of the architecture and reference design, providing technologies that meet the project requirements and assisting in the installation and configuration of those technologies. The practice guide highlights the approach used to develop the NCCoE reference solution. Elements include risk assessment and analysis, logical design, build development, test and evaluation, and security control

mapping. This guide is intended to provide practical guidance to any organization interested in implementing a solution for recovery from a cybersecurity event.

3.1 Audience

This guide is intended for individuals responsible for implementing security solutions in organizations' IT support activities. Current IT systems, particularly in the private sector, often lack integrity protection for domain name services and electronic mail. The platforms demonstrated by this project, and the implementation information provided in these practice guides, permit integration of products to implement a data recovery system. The technical components will appeal to system administrators, IT managers, IT security managers, and others directly involved in the secure and safe operation of the business IT networks.

3.2 Scope

The guide provides practical, real-world guidance on developing and implementing a DI solution consistent with the principles in the *NIST Framework for Improving Critical Infrastructure Cybersecurity Volume 1* [1], specifically the core function of recover. Recover emphasizes developing and implementing the appropriate activities to maintain plans for resilience and to restore any capabilities or services that were impaired by a cybersecurity event to a last known good state. Examples of outcomes within this function include recovery planning, improvements, and communication.

3.3 Assumptions

This project is guided by the following assumptions:

- The solution was developed in a lab environment. The environment is based on a typical organization's IT enterprise. It does not reflect the complexity of a production environment.
- An organization has access to the skill sets and resources required to implement a data recovery solution.
- A DI event has taken place and been detected. This guide does not address the actual detection function.

3.4 Risk Assessment

NIST SP 800-30 Rev. 1: Guide for Conducting Risk Assessments [3] states that the definition of risk is "a measure of the extent to which an entity is threatened by a potential circumstance or event, and typically a function of: (i) the adverse impacts that would arise if the circumstance or event occurs; and (ii) the likelihood of occurrence." The NCCoE recommends that any discussion of risk management, particularly at the enterprise level, begin with a comprehensive review of *NIST 800-37: A Guide for Applying the Risk Management Framework to Federal Information Systems* [4]. The framework proved

invaluable in giving us a baseline to assess risks, from which we developed the required security controls of the reference design and this guide.

We performed two types of risk assessment:

- Initial analysis of the risk factors that were discussed with financial, retail, and hospitality institutions. This analysis led to the creation of the DI project and the desired security posture. See *NIST IR 8050 Executive Technical Workshop* [2] for additional participant information.
- Analysis of how to secure the components within the solution and minimize any vulnerabilities they might introduce. See Section 6, Security Characteristics Analysis.

3.4.1 Assessing Risk Posture

Using the guidance in NIST's series of publications concerning risk, we worked with financial institutions and the Financial Sector Information Sharing and Analysis Center to identify the most compelling risk factors encountered by this business group. We participated in conferences and met with members of the financial sector to define the main security risks to business operations. These discussions resulted in the identification of an area of concern—the inability to recover from DI attacks. We then identified the core operational risks, as various methods exist that all lead to sustaining a DI compromise. These risks lead to two tactical risk factors:

- systems incapacitated
- DI impacted

These discussions also gave us an understanding of strategic risks for organizations with respect to DI. *NIST SP 800-39: Managing Information Security Risk* [5] focuses particularly on the business aspect of risk, namely at the enterprise level. This understanding is essential for any further risk analysis, risk response/mitigation, and risk monitoring activities. The following is a summary of the strategic risk areas we identified and their mitigations:

- Impact on system function – ensuring the availability of accurate data or sustaining an acceptable level of DI reduces the risk of systems' availability being compromised.
- Cost of implementation – implementing DI once and using it across all systems may reduce both system restoration and system continuity costs.
- Compliance with existing industry standards – contributes to the industry requirement to maintain a continuity of operations plan.
- Maintenance of reputation and public image – helps reduce level of impact, in turn helping to maintain image.
- Increased focus on DI – includes not just loss of confidentiality but also harm from unauthorized alteration of data (per NIST IR 8050 [2]).

We subsequently translated the risk factors identified to security functions and subcategories within the NIST CSF. In Table 3-1 we mapped the categories to NIST's *SP 800-53 Rev. 4* [7] controls and International Electrotechnical Commission/International Organization for Standardization (IEC/ISO) controls for additional guidance.

3.4.2 Security Control Map

As explained in Section 3.4.1, we identified the CSF security functions and subcategories that we wanted the reference design to support through a risk analysis process. This was a critical first step in designing the reference design and example implementation to mitigate the risk factors. Table 3-1 lists the addressed CSF functions and subcategories and maps them to relevant NIST standards, industry standards, and controls and best practices. The references provide solution validation points in that they list specific security capabilities that a solution addressing the CSF subcategories would be expected to exhibit. Organizations can use Table 3-1 to identify the CSF subcategories and NIST 800-53 controls that they are interested in addressing.

Note: Not all the CSF subcategories guidance can be implemented using technology. Any organization executing a DI solution would need to adopt processes and organizational policies that support the reference design. For example, some of the subcategories within the CSF function "Identify" are processes and policies that should be developed prior to implementing recommendations.

Table 3-1 Data Integrity Reference Design CSF Core Components Map

Cybersecurity Framework (CSF) v1.1				Standards & Best Practices
Function	**Category**	**Subcategory**	**SP800-53R4**	**ISO/IEC 27001:2013**
PROTECT (PR)		PR.DS-1: Data-at-rest is protected	SC-28	A.8.2.3
		PR.DS-6: Integrity checking mechanisms are used to verify software, firmware, and information integrity	SI-7	A.12.2.1, A.12.5.1, A.14.1.2, A.14.1.3

Cybersecurity Framework (CSF) v1.1				Standards & Best Practices
Function	**Category**	**Subcategory**	**SP800-53R4**	**ISO/IEC 27001:2013**
	Information Protection Processes and Procedures (PR.IP)	PR.IP-3: Configuration change control processes are in place	CM-3, CM-4, SA-10	A.12.1.2, A.12.5.1, A.12.6.2, A.14.2.2, A.14.2.3, A.14.2.4, A.14.2.7
		PR.IP-4: Backups of information are conducted, maintained, and tested periodically	CP-4, CP-6, CP-9	A.11.1.4, A.12.3.1, A.17.1.2, A.17.1.3, A.17.2.1 A. 18.1.3
		PR.IP-9: Response plans (Incident Response and Business Continuity) and recovery plans (Incident Recovery and Disaster Recovery) are in place and managed	CP-2, IR-8	A.16.1.1, A.17.1.1, A.17.1.2, A.17.2.1
	Protective Technology (PR.PT)	PR.PT-1: Audit/log records are determined, documented, implemented, and reviewed in accordance with policy	AU Family IR-5, IR-6	A.6.1.3, A.16.1.2, A.12.4.1, A.12.4.2, A.12.4.3, A.12.4.4, A.12.7.1

Cybersecurity Framework (CSF) v1.1				Standards & Best Practices
Function	**Category**	**Subcategory**	**SP800-53R4**	**ISO/IEC 27001:2013**
DETECT (DE)	Anomalies and Events (DE.AE)	DE.AE-4: Impact of events is determined	CP-2, IR-4, RA-3, SI -4	A.6.1.1, A.17.1.1, A.17.2.1, A.16.1.4, A.16.1.5, A.16.1.6, A.12.6.1
	Security Continuous Monitoring (DE.CM)	DE.CM-1: The network is monitored to detect potential cybersecurity events	AC-2, AU-12, CA-7, CM-3, SC-5, SC-7, SI-4	A.9.2.1, A.9.2.2, A.9.2.3, A.9.2.5, A.9.2.6, A.12.4.1, A.12.4.3, A.12.1.2, A.14.2.2, A.14.2.3, A.14.2.4, A.13.1.1, A.13.1.3, A.13.2.1, A.14.1.3
		DE.CM-3: Personnel activity is monitored to detect potential cybersecurity events	AC-2, AU-12, AU-13, CA-7, CM-10, CM-11	A.9.2.1, A.9.2.2, A.9.2.3, A.9.2.5, A.9.2.6, A.12.4.1, A.12.4.3, A.18.1.2, A.12.5.1, A.12.6.2s

3.5 Technologies

Table 3-2 lists all the technologies used in this project and provides a mapping between the generic application term, the specific product used, and the security control(s) that the product provides. Refer to Table 3-1 for an explanation of the CSF subcategory codes. This table describes only the product capabilities used in our example solution. Many of the products have additional security capabilities that were not used for our purposes.

379 **Table 3-2 Products and Technologies**

Component	Specific Product	Function	CSF Subcategories
Corruption Testing	ArcSight Enterprise Security Manager (ESM) v6.9.1 Tripwire Enterprise v8.5 Tripwire Log Center Manager v7.2.4.80	• provides monitoring for changes to data on a system • provides logs, detection, and reporting, in the event of changes to data on a system • provides audit capabilities for database metadata and content modifications • provides file hashing and integrity testing independent of file type (can include software files) • provides notifications for changes to configuration • provides file monitoring for cybersecurity events • provides analytic capabilities to determine the impact of integrity events	PR.DS-6, PR.PT-1, DE.AE-4
Secure Storage	Spectrum Protect and Backup and Replication v8.1.0 WORMdisk v151228	• provides write-once read-many file disk storage for secure backups of integrity information • provides immutability of backups • creates encrypted backups	PR.DS-1, PR.IP-4
Logging	ArcSight Enterprise Security Manager (ESM) v6.9.1	• provides auditing and logging capabilities configurable to corporate policy • provides logging of some user activity of monitored systems	PR.PT-1, DE.AE-4, DE.CM-1, DE.CM-3

Component	Specific Product	Function	CSF Subcategories
	Tripwire Enterprise v8.5 Tripwire Log Center Manager v7.2.4.80	• provides network information about certain cybersecurity events • correlates logs of cybersecurity events with user information • provides logs of database activity and database backup operations • provides analysis capabilities for log data • provides analysis capabilities for finding anomalies in user activity • provides automation for logging • provides logs of database activity and database backup operations	
Backup Capability	Spectrum Protect and Backup and Replication v8.1.0 WORMdisk v151228	• provides backup and restoration capabilities for systems • provides backup and restore capabilities for configuration files • provides immutable storage • performs periodic backups of information	PR.DS-1, PR.IP-3, PR.IP-4, PR.IP-9
Virtual Infrastructure	Veeam Availability Suite 9.5	• provides backup and restoration capabilities for virtual systems • provides ability to encrypt backups • provides logs for backup and restore operations	PR.DS-1, PR.IP-4, PR.PT-1

380

4 Architecture

Data integrity involves the recovery of data after a ransomware or other destructive attack with the validation that the recovered data is the last known good. This section presents a high-level architecture and reference design for implementing such a solution.

4.1 Architecture Description

4.1.1 High-Level Architecture

The DI solution is designed to address the security functions and subcategories described in Table 3-1 and is composed of the capabilities illustrated in Figure 4-1.

Figure 4-1 DI High-Level Architecture

1. Secure Storage provides the capability to store data with additional data protection measures, such as Write Once Read Many (WORM) technologies or data encryption.
2. Logging stores and reports all the log files produced by the components within the enterprise.
3. Virtual Infrastructure provides virtualized capabilities, including backup capabilities for the virtual infrastructure.
4. Corruption Testing provides capabilities for testing file corruption and provides notification or logs of violations against specified policies.
5. Backup Capability establishes a capability for components within the enterprise that are not a part of the virtual infrastructure to produce a backup.

These capabilities work together to provide the recover function for DI. The secure storage is the ability to store file-such as backups, gold images, or configurations files, in a format that cannot be corrupted, since files cannot be altered or changed while in storage. The logging capability works in conjunction with the corruption testing. The corruption testing capability describes the event(s) when the attack occurs and the damage caused. Since the corruption testing describes when the event occurred, these details can be used to investigate the logs to correlate all events relative to the attack across all items

406 that report log files. After the last known good is determined via the logs and corruption testing, the
407 backup capability for either the enterprise or the virtual infrastructure is employed. A backup capability
408 is the ability to restore to the point prior to the DI event. The backup capability is supplemented by built-
409 in backup and rollback capabilities of the database services.

410 The following components of the high-level architecture are not addressed in this guide: enterprise
411 components (e.g., virtual machines, mail servers, active directory, file sharing capabilities), installation
412 and configurations, file corruption testing policies, and event detection.

413 4.1.2 Reference Design

414 The reference design addresses the DI architecture in conjunction with its interactions with a
415 representation of a basic enterprise.

416 **Figure 4-2 DI Reference Design**

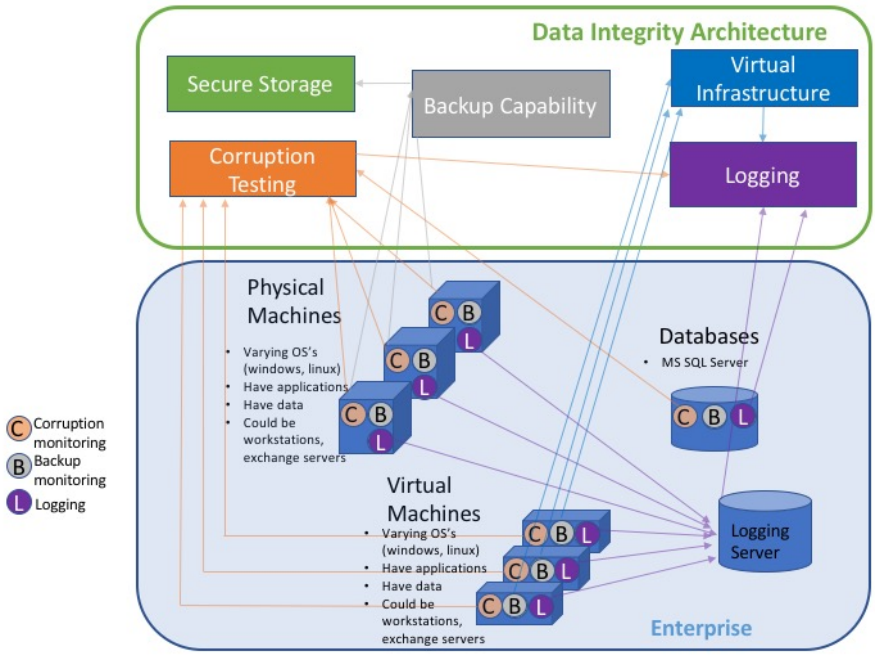

417

418 Solid lines represent the communication of information between components within the enterprise,
419 from the enterprise to the DI architecture, or between components within the DI architecture. The lines
420 are color coded to correspond with the capability provided by the DI architecture.

421 The Secure Storage component provides a capability to store the most critical files for an enterprise.
422 These would include backup data, configuration files, and golden images. Additional measures need to
423 be applied to provide increased security to these files so they are not subject to attacks or corruption.

The Corruption Testing component provides the ability to test, understand, and measure the attack that occurred to files and components within the enterprise. This testing is essential to identify the last known good for the DI recovery process. For these measures to be applicable to an enterprise, appropriate triggers need to be defined and developed within the capability that look for specific events. For example, it may be very normal for end users to have encrypted files they develop during operational hours. But if every file on the end user's workstation begins to be encrypted, or an encryption begins to happen on the end user machine at hours outside of normal operational hours, these could be identifiable actions noted in the log files indicating a ransomware attack. For an enterprise, these triggers need to be defined appropriately and thoroughly to have a successful Corruption Testing capability.

The Backup Capability component supports the ability to back up each component within the enterprise as well as perform a restore that uses backup data. The configuration of this component needs to align with the tempo of the enterprise. For example, if an enterprise is performing thousands of transactions per hour per day, then a backup solution that only performs a backup once a day would not adequately provide for the enterprise. This type of configuration would allow for a potentially large data loss. If backups occur every morning and a loss of DI happened at the end of the day, then a full day's worth of transactions would be lost. The decision on what the correct configuration is determined by an organization's risk tolerance. More information pertaining to this decision can be found in Section 5.1.1.3.

The Virtual Infrastructure component straddles the line between being part of the enterprise and part of the DI architecture. It provides virtual capabilities to the enterprise as well as backup and restoration capabilities to support the DI architecture. The backup and restoration capabilities are for the virtual infrastructure itself. For data that is produced on individual virtual machines (VMs), either the VM infrastructure can provide the file-level restoration or the backup component can provide this capability. If the VM infrastructure cannot provide its own backup and restoration, then the requirements for that are levied on the backup component.

Logging from each component and sorting the logs together is imperative to understanding the ramifications of the attack across the enterprise. File, system, and configuration changes and modifications need to be logged, reported, and stored in one repository where events can be identified and understood.

Databases are necessary to support everyday operations of the enterprise architecture and to assist in backup and recovery. The chosen database software should have built-in backup and rollback methods enabled, although commercial solutions for the backup and recovery of databases exist. Often, these commercial solutions use the internal database backup/recovery capabilities. These capabilities are tied into the security architecture, as demonstrated in Section 5.1.6.2. Consult the Backup Capability paragraph above for guidance on the regularity of backups. The regularity of database backups determines the effectiveness of data recovery efforts.

5 Example Implementation

The example implementation is constructed on the NCCoE lab's infrastructure, which consists of a VMware vSphere virtualization operating environment. We used network attached storage and virtual switches, as well as internet access, to interconnect the solution components. The lab network is not connected to the NIST enterprise network. Table 5-1 lists (alphabetically) the software and hardware components we used, as well as the specific function each component.

Table 5-1 Example Implementation Component List

Product Vendor	Component Name	Function
GreenTec	WORMdisk	Secure, immutable hardware
Hewlett Packard Enterprise (HPE)	ArcSight ESM	Log analysis, correlation, management, and reporting
IBM	Spectrum Protect	File-level, disk-level, and system-level backup and recovery
Tripwire	Enterprise and Log Center	File integrity monitoring and database metadata integrity monitoring
Veeam	Availability Suite	VM backup and restore

The architecture depicted in Figure 5-1 describes a solution built around several typical infrastructure components: a Microsoft Exchange server, a Microsoft SharePoint server, a Microsoft Structured Query Language (MS SQL) server, a Microsoft Hyper-V server, and a Microsoft Active Directory server that also runs Microsoft Domain Name System service, as well as an array of client machines, primarily running Windows 10 and Ubuntu 16.04.

The solution consists of several products to comprise an enterprise DI solution.

Organizations should have backup capability that can be used to back up files, disks, and systems. Tools that provide backup capability may also provide capabilities to back up databases or email servers. These tools should include management capabilities for backups that provide configuration options such as when and how data should be backed up. IBM Spectrum Protect provides backup capability in this build. Clients are installed on all machines that need backup and restore capabilities. Furthermore, IBM Spectrum Protect uses incremental backups; essentially, this means that it stores an initial full backup of a user's system. After this initial backup, additional backups are performed only after changes occur in data.

Secure storage is important for protecting backups and other forms of data in an enterprise DI solution. Secure storage involves write-protected or write-controlled devices, which prevent data from being modified or deleted. By integrating backup infrastructure with these disks, it is possible to permanently

preserve backups and protect them from harmful malware and accidental deletion. GreenTec WORMdisks are a secure storage solution that protects data on a firmware level. WORMdisks come with software to lock disks or portions of disks permanently or temporarily. Once WORMdisks are locked, they are immutable and any data on the disk is read-only. Implementation instructions are included for backing up directly to GreenTec WORMdisks using IBM Spectrum Protect, as well as instructions for copying backup data from IBM Spectrum Protect to a WORMdisk. Other files stored on these disks can be copied over using the operating system's usual methods. WORMdisks are transparent to the operating system in terms of use, so they function as regular storage drives until they are locked.

Corruption testing involves periodic or manual testing of files for modifications, deletions, additions, or other potential DI events. Tools that provide corruption testing may also test other systems, such as databases or mail servers. Tripwire Enterprise provides corruption testing for this build. By using individual agents installed on client machines, Tripwire Enterprise generates file integrity information for a set of specified files and folders. Tripwire Enterprise can also generate file integrity information for database metadata, allowing administrators to track changes made to database structure. It stores this metadata in a database. For simplicity, we use the MS SQL server to store the file integrity information, but this could be done in a separate database for processing efficiency. Tripwire Enterprise forwards logs that it generates to Tripwire Log Center. Tripwire Log Center allows for filtering and processing of Tripwire Enterprise logs as well as the ability to integrate with other log collection tools.

Many organizations have virtual infrastructure that allows them to manage the distribution of VMs across their enterprise. When implementing a DI solution, the virtual infrastructure should include the ability to granularly backup and restore VMs. Veeam Backup and Replication is a tool that can integrate with Hyper-V and VMware to jointly comprise the virtual infrastructure of our build. Veeam Backup and Replication can provide granular backup and restore capabilities. It can perform restores of entire VMs as well as restores on individual files in virtualized environments. Veeam Backup and Replication is server based and can be applied to Hyper-V machines that run on various systems across the enterprise.

Logging is another important piece of a DI solution. The collection of logs from various sources is useful in identifying the root cause of DI events, whether they are caused by accident or by malicious insiders or software. Furthermore, logs aid in identifying the time of the last known good and inform decisions regarding restoration. In this build, HPE ArcSight ESM is used to collect logs from various sources. Included in the architecture is an HPE ArcSight Connector server. Through Active Directory, the connector server acquires system and security logs from all Windows endpoints in the domain. These logs are then forwarded to HPE ArcSight ESM. Implementation instructions are included for other, non-default sources. HPE ArcSight ESM can log MS SQL queries and collect Hyper-V application logs, Veeam application logs, and Ubuntu syslogs, and provides instructions for each. In the case of Hyper-V application logs and Veeam application logs, we provide sample custom parsers for forwarding some events to HPE ArcSight ESM (see Volume 3). Additionally, ESM integrates with Tripwire Log Center to provide log collection for all file integrity monitoring logs generated by Tripwire Enterprise. HPE ArcSight ESM can sort, filter, and audit logs from all its sources. The information gathered from these logs should

provide system administrators the context they need to determine how to fully remediate systems affected by destructive malware.

Figure 5-1 Example Implementation Architecture

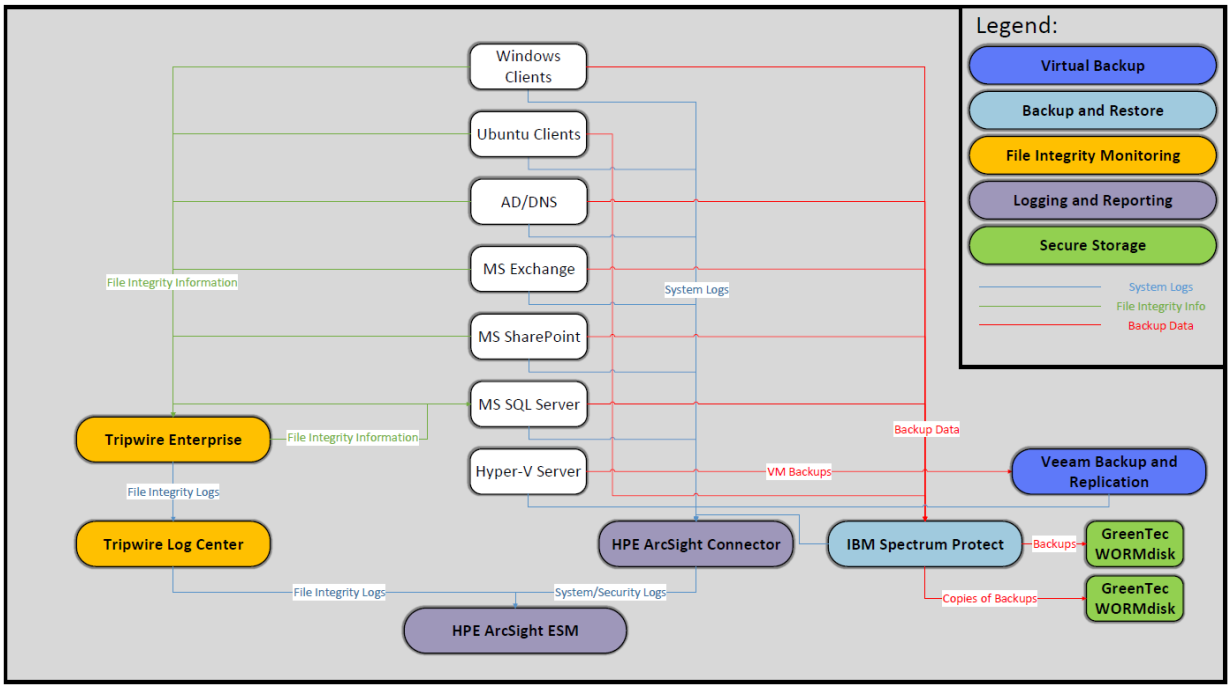

5.1 Use Cases

5.1.1 Ransomware

5.1.1.1 Scenario

A malicious piece of software run by the user encrypts the entire documents folder. This renders files unusable and pictures unable to be viewed, and users will only be able to see encrypted text should they attempt to open any of the files in a text editor. Though the software's scope is limited to the documents folder, the approach could be more widely applied to encrypt other folders and even system files, resulting in an attack on the availability of systems and data alike.

5.1.1.2 Resolution

This use case is resolved using a combination of several tools. The corruption testing component (Tripwire Enterprise) is used to detect changes in the file systems of various selected machines, specifically when files are modified or overwritten. The corruption testing component provides context

for these events, such as a time stamp, the user responsible, the affected files, and the program that modified the file (if applicable).

The logging component (HPE ArcSight ESM) collects logs from various sources for analysis and reporting. Logs are forwarded from the corruption testing component for analysis by a system administrator. The logging component provides search, filtering, and correlation capabilities for auditing, allowing enterprises to manage the quantity of logs generated by the corruption testing component and other sources.

These two components work together to provide information about the files encrypted by the ransomware tool: the name of the program that encrypted the files, which files were affected, when they were affected, and which user ran the program. This information aids in removing the ransomware from the system and contributes to the identification of the last known good. However, it does not actually restore the availability of the user's files. The backup capability component (IBM Spectrum Protect) is used to restore encrypted files.

5.1.1.3 Other Considerations

In the event of a system failure caused by ransomware, it is important to note that recovery requires the installation of the IBM Spectrum Protect client (if IBM Spectrum Protect is used as the backup capability). If a system failed due to ransomware and cannot be rebooted, this client may not be immediately accessible. Restoration would require the reinstallation of the operating system and then installation of the IBM Spectrum Protect client. The client could then restore all files, including system files, to their previous state. Products exist that work with IBM Spectrum Protect to automate and accelerate this process.

Also, there is a trade-off between the frequency of backups and the amount of data loss an enterprise will experience. More frequent backups require more resources, both in work performed by the client and space required on the server. More frequent backups, however, provide more granularity in recovery capabilities. This can be managed by backing up active files more frequently and dormant files less frequently. An active file will lose more data during recovery because the restoration is to a point in time and will not reflect recent changes to the file.

Another caveat of more frequent (i.e., automated) backups is that if a backup is taken after a ransomware attack, the backup infrastructure will retain backups of the encrypted data. Though this is undesirable, it is still possible to restore to previous versions. This scenario highlights the importance of file monitoring capabilities, which can guide users to restoring to the correct backup.

5.1.2 File Modification and Deletion

5.1.2.1 Scenario

A malicious piece of software is downloaded from a phishing website and run by the user. The software recursively modifies files in the directory in which it is running. It removes and replaces pieces of text files, such as numbers and common English words, sometimes removing entire lines of text. It also deletes any file it doesn't recognize as text, such as pictures, videos, and music files. This results in potentially detrimental data loss. Furthermore, since files are deleted and not just encrypted, recovery is impossible without a backup infrastructure in place. There is no option to decrypt files that were deleted from the system, so compensating the creators of the malicious software for data recovery is not an option.

5.1.2.2 Resolution

Though this use case is more destructive than ransomware, the same tools are used to recover from it. The corruption testing component (Tripwire Enterprise) is used to test sensitive files and folders, and reports information such as the time, user, and the name of the malicious software that deleted and modified the now corrupted files. Even though files are missing and not just encrypted, their deletion will still be reported.

The logs generated by the corruption testing component are forwarded to the logging component (HPE ArcSight ESM) for collection and processing by a system administrator. The administrator can use the information to determine how to respond to the event—how to remove the malicious software, how to prevent it from spreading, and which files to restore. The combination of logging in concert with corruption testing provides the ability to identify the last known good.

The backup capability (IBM Spectrum Protect) is used to restore modified, corrupted, and deleted files. Even though files are missing from the user's system, they are still present in the backup capability component, and the user need only choose which backup version to restore to.

5.1.2.3 Other Considerations

Please see Section 5.1.1.3 for a discussion of tradeoffs between the frequency of backups, resources required, and restoration granularity, as they are applicable to this use case.

Again, if a backup is taken after malicious software runs but before recovery, the corrupted data will be retained by the backup infrastructure. However, it will still be possible to restore to an older version of the data with IBM Spectrum Protect (if IBM Spectrum Protect is used). IBM Spectrum Protect will not back up deleted files, however, so in the event of file deletion, the last backup taken should be sufficient for recovery, unless the user has a specific reason to recover from an earlier version.

5.1.3 VM Deletion

5.1.3.1 Scenario

A user accidentally deleted a VM in Hyper-V. In this use case, it is assumed that the user has access to the VM. Although the deletion may not set off any red flags by detection systems since a privileged user deleted the machine, it is still undesired. Since VMs can be used for several purposes—such as access to software unavailable on the host operating system (OS), emulation of infrastructure before deployment, or simply storing files for use in the user's preferred OS—the deletion of a VM can cause significant data loss and disruption in work flow.

5.1.3.2 Resolution

The VM deletion is resolved using a combination of the logging component (HPE ArcSight ESM) and the virtual infrastructure (Veeam Backup and Restore, Hyper-V). This use case deals specifically with an accidental deletion by a benign user. Because of this, logs pertaining to the deletion are likely unnecessary for recovery. However, other use cases may require logs, especially in the event of a malicious VM deletion. Therefore, our resolution includes a method for integrating the selected virtual infrastructure tools and logging component. The integration allows for the collection of logs regarding the deletion of the VM as well as logs pertaining to the restoration of the VM once complete. The virtual infrastructure is used to restore the entire deleted VM.

5.1.3.3 Other Considerations

The chosen virtual infrastructure components (Veeam Backup and Restore, Hyper-V) allow for more granular recovery–files on the guest OS can be recovered, not just the entire VM. This extends the user's restoration capabilities in events where data corruption happens within the VM. However, it is unlikely that file change logs will be forwarded to the logging component (HPE ArcSight ESM), meaning that such recovery capabilities do not meet all the requirements of this reference design.

5.1.4 Active Directory Permission Change

5.1.4.1 Scenario

A malicious insider creates backdoors into a Microsoft Exchange server. Since the culprit is an insider, he or she is assumed to be privileged. The backdoor accounts have administrator privileges and can make changes to various settings in the Exchange infrastructure. This results in potential data leaks, which could involve forwarding emails from all users to an off-site account.

5.1.4.2 Resolution

This use case is resolved primarily using the logging component (HPE ArcSight ESM) and the built-in Microsoft Windows server recovery capabilities. Since system and security logs are reported to the

logging component, administrators will be able to find which user created the accounts, the names of all the accounts created, when they were created, and the account activities. The administrator could choose to delete the accounts manually, but Windows includes a method for restoring the system state. Since restoring the system state is more complicated in later Windows server versions, the chosen backup capability (IBM Spectrum Protect) is not used for the restoration. As stated in the product documentation, the preferred method for recovering the system state is through the Microsoft Windows System State restoration process.

This restore is performed on the Active Directory server (as opposed to the Microsoft Exchange server) since the accounts, though created from the Exchange server, are stored on the Active Directory server.

5.1.4.3 Other Considerations

IBM Spectrum Protect recommends using the Microsoft Windows System State backup and recovery tool for later Windows versions.

5.1.5 Database Transactions

5.1.5.1 Scenario

A malicious or careless insider changes database data that is necessary for enterprise operations. The user is assumed to be privileged. Through the course of interacting with the database, the user executes a query that inserts, deletes, or modifies data in a way that harms enterprise operations.

5.1.5.2 Resolution

The event is detected with the logging capability (HPE ArcSight ESM). Database integrity is restored through a system of transactional rollbacks. Since the logging capability includes database query log collection, administrators will be able to find which users modified the database, and what queries were run. Given this information, administrators can determine the harmful queries and when the database was in its desired state. Transactional rollbacks are then used to restore the database to the last known good state.

5.1.5.3 Other Considerations

Restoration need not be conducted on the database server, depending on the method of rollbacks employed. The database modification can be conducted on any machine.

Transactional rollbacks require that queries be explicitly executed within “transactions.” During the restoration process, a transactional ID is specified to restore to. An enterprise can choose to force queries to use transactions through the implementation of a proxy between all potential endpoints and the database. Through this precise processing of queries, granular restoration can be achieved, though potentially at cost to efficiency.

5.1.6 Database Metadata Modification

5.1.6.1 Scenario

A malicious or careless insider changes the metadata of the system's main database. The user is assumed to be privileged. Through the course of interacting with the database, the user executes a query that changes the name of a key table. This results in a loss of functionality of the database for any queries that wish to use that table.

5.1.6.2 Resolution

This use case is resolved through database restoration capabilities—in this case, inherent to the database. Both the corruption testing component (Tripwire Enterprise) and the logging component (HPE ArcSight ESM) are used to detect the event. Through these components, administrators will be able to find which users modified the database. It is possible to manually revert the changes, but the built-in database backup and restoration capabilities can also be used to fix the metadata.

Regardless of where the database modification query was run, recovery occurs on the database server to the last known good.

5.1.6.3 Other Considerations

Backup scheduling tied to the database is separate from the backup capability (IBM Spectrum Protect). If tools are used that require separate database backup procedures, security policies and backup schedules should be designed to accommodate this fact.

Note: The use of backups to restore databases that have had adverse changes to their metadata may result in the loss of all data since the backup was taken. Reversing the changes manually is more time-consuming but more precise.

6 Security Characteristics Analysis

This evaluation focuses on the security of the reference design itself. In addition, it seeks to understand the security benefits and drawbacks of the example solution.

6.1 Assumptions and Limitations

The security characteristic evaluation has several limitations:

- It is not a comprehensive test of all security components, nor is it a red team exercise.
- It cannot identify all weaknesses.

- It does not include the lab infrastructure. It is assumed that devices are hardened. Testing these devices would reveal only weaknesses in implementation that would not be relevant to those adopting this reference architecture.

6.2 Analysis of the Reference Design's Support for CSF Subcategories

Table 3-2 lists the reference design functions and the security characteristics, along with products that we used to instantiate each capability. The focus of the security evaluation is not on these specific products but on the CSF subcategories, because, in theory, any number of commercially available products could be substituted to provide the CSF support represented by a given reference design capability.

This section discusses how the reference design supports each of the CSF subcategories listed in Table 3-1. Using the CSF subcategories as a basis for organizing our analysis allowed us to systematically consider how well the reference design supports specific security activities and provides structure to our security analysis.

6.2.1 PR.IP-3: Configuration Change Control Processes Are in Place

The reference design protects the configuration from change and detects changes in the configuration using secure hardware and file integrity monitoring. It does not include processes for change control, however, which the adopting organization should implement.

6.2.2 PR. IP-4: Backups of Information Are Conducted, Maintained, and Tested Periodically

The reference design includes capabilities for creating backups of information from various sources:

- file systems
- disks
- virtualized environments
- databases

It also describes scheduling capabilities for each of these backup targets, allowing for periodic backups as well as manual backups. The design provides the capability to test and maintain backups, but planning schedules, maintenance, and testing of backups are left to the adopting organization.

By adopting this reference design, organizations gain the capability to conduct, maintain, and test backups, and in doing so, the organizations will support the technical requirements of CSF subcategory PR.IP-4.

6.2.3 PR.DS-1: Data-at-Rest Is Protected

The reference design supports the protection of data-at-rest through:

- secure hardware as protection against data corruption
- encryption of backups as protection against unauthorized access

Through these combined capabilities, the reference design can protect data-at-rest from both unauthorized reads and writes. This protection only applies to data that is stored using the capability of the reference design. Utilization of the reference design is necessary for data protection; implementation alone is not sufficient.

By adopting this reference design, organizations gain the capability to protect data-at-rest, and in doing so, the organizations will support the technical requirements of CSF subcategory PR.DS-1.

6.2.4 PR.DS-6: Integrity Checking Mechanisms Are Used to Verify Software, Firmware, and Information Integrity

The reference design supports integrity checking for various types of data, including:

- files stored in file systems
- database metadata
- logs
- software

Firmware that is stored on special hardware may be out of the scope of the design. It should be possible to monitor firmware stored as files; however, this reference design does not include firmware or software integrity verification against online resources.

By adopting this reference design, organizations gain the capability to monitor file integrity within their system. This partially supports the technical requirements of CSF subcategory PR.DS-6, but the verification of integrity for firmware and software against verified sources is out of scope.

6.2.5 PR.PT-1: Audit/Log Records Are Determined, Documented, Implemented, and Reviewed in Accordance with Policy

The reference design supports auditing, log collection, log analysis, and log correlation. It includes mechanisms for collecting logs from:

- Microsoft event logs
- Windows application logs
- Linux system logs

- file integrity logs
- custom log sources
- database query history

Logs are aggregated into a single interface, which allows for searching, correlating, and analyzing logs from across an enterprise. Reviewing these logs is left to the individual organization.

By adopting this reference design, organizations gain the technical capability to aggregate, correlate, and analyze logs as well as perform audits across an enterprise. In doing so, the organizations will support the technical requirements of CSF subcategory PR.PT-1.

6.2.6 DE.CM-3: Personnel Activity Is Monitored to Detect Potential Cybersecurity Events

The reference design supports log collection for various activities across an enterprise, including:

- file creation, deletion, modification, and renaming
- account creation, deletion, and modification
- database queries and other activity

These collected logs, where possible, have users and programs associated with them. The design does not support active monitoring of user activity or monitoring of network activity. However, logs are provided for relevant activities, so that informed decisions can be made when an organization decides how to recover from destructive malware.

By adopting this reference design, organizations will gain the technical capability to review some personnel activity after a cybersecurity event has occurred, and in doing so, partially support the technical requirements of CSF subcategory DE.CM-3.

6.2.7 DE.CM-1: The Network Is Monitored to Detect Potential Cybersecurity Events

The reference design supports the monitoring of some network activity in the enterprise. Network information is correlated with all logged cybersecurity events to determine:

- Source Internet Protocol (IP) of event (if applicable)
- Destination IP of event (if applicable)
- Port (if applicable)

Though these collected logs have network information associated with them, network activity is not directly monitored for anomalies. Since the focus of this project is recovery, the reference design supports enough network information to recover from a cybersecurity event, but will not attempt to detect cybersecurity events based on network traffic or packet analysis.

By adopting this reference design, organizations will gain the technical capability to associate DI events with network information, and in doing so, will partially support the technical requirements of CSF subcategory DE.CM-1.

6.2.8 DE.CM-2: The Physical Environment Is Monitored to Detect Potential Cybersecurity Events

The reference design supports the monitoring of physical machines in the enterprise through the real-time monitoring of:

- file integrity
- database metadata integrity
- database queries

This reference design does not include monitoring for physical cybersecurity events, such as the insertion of potentially malicious flash drives.

By adopting this reference design, organizations will only partially gain the technical capability required to fully monitor the physical environment, and in doing so, partially support the technical requirements of CSF subcategory DE.CM-2.

6.2.9 PR.IP-9: Response Plans and Recovery Plans Are in Place and Managed

The reference design supports notification after a DI event as well as the infrastructure required for recovery, including:

- logs for analysis and auditing events after they happen
- backup and restore capabilities for successful recovery

The design supports the technical requirements of a recovery plan; however, the details of the plan should be put in place by the adopting organizations.

By adopting this reference design, organizations will gain the technical capability required to recover from a DI event, and in doing so, support the technical requirements of CSF subcategory PR.IP-9.

6.2.10 DE.AE-4: Impact of Events Is Determined

The reference design supports an infrastructure to determine the scope of DI events as well as create plans of action for remediation. This infrastructure includes:

- logs that identify impacted files and systems
- auditing to determine responsible parties after an event occurs

The design provides the forensic ability to determine affected systems and responsible parties but does not act on this information without human intervention. Adopting organizations should create plans to use this information for remediation.

By adopting the design, organizations will only partially gain the technical capability required to determine the impact of events, and in doing so, partially support the technical requirements of CSF subcategory DE.AE-4.

6.3 Security of the Reference Design

The list of reference design capabilities in Table 3-2 focuses on the capabilities needed to ensure the integrity of system data. Table 3-2 does not focus on capabilities that are needed to manage and secure the reference design. However, the reference design itself must be managed and secured. To this end, this security evaluation focuses on the security of the reference design itself.

Measures implemented to protect the reference design from outside attack include:

- isolating certain capabilities on separate subnetworks protected by firewalls
- Implementing a management network to isolate log and management traffic from the production (business operations) networks
- securing critical user access information and logs to protect them from unauthorized insertion, modification, or deletion
- logging all privileged user access activities
- using encryption and integrity protection of user access information and logs while this information is in transit between capabilities

Table 6-1, Capabilities for Managing and Securing the DI Reference Design, describes the security protections each capability provides and lists the corresponding products that were used to instantiate each capability. The security evaluation focuses on the capabilities rather than the products. The NCCoE is not assessing or certifying the security of the products included in the example implementation. We assume that the enterprise already deploys network security capabilities such as firewalls and intrusion detection devices that are configured per best practices. The focus here is on securing capabilities introduced by the reference design and minimizing their exposure to threats.

6.3.1 Deployment Recommendations

When deploying the reference design in an operational environment, organizations should follow security best practices to address potential vulnerabilities and ensure that all solution assumptions are valid to minimize any risk to the production network. Organizations leveraging the reference design should adhere to the following list of recommended best practices that are designed to reduce risk. Note that the laboratory instantiation of the reference design did not implement every security

recommendation. Organizations should not, however, consider this list to be comprehensive; merely following this list will not guarantee a secure environment. Organizations must also take into consideration items such as user access controls, continuity of operations planning, and environmental elements that are not addressed in this document. Planning for design deployment gives an organization the opportunity to go back and audit the information in its system and get a more global, correlated, and disambiguated view of the DI controls that are in effect.

6.3.1.1 Patch, Harden, Scan, and Test [6]

- Keep OSs up-to-date by patching, version control, and monitoring indicators of compromise (e.g., performing virus and malware detection as well as keeping anti-virus signatures up-to-date).
- Harden all capabilities by deploying on securely configured OSs that use long and complex passwords and are configured per best practices.
- Scan OSs for vulnerabilities.
- Test individual capabilities to ensure that they provide the expected CSF subcategory support and that they do not introduce unintended vulnerabilities.
- Evaluate reference design implementations before going operational with them.

6.3.1.2 Other Security Best Practices [7]

- Install, configure, and use each capability of the reference design per the security guidance provided by the capability vendor.
- Change the default password when installing software.
- Identify and understand which predefined administrative and other accounts each capability comes with by default to eliminate any inadvertent backdoors into these capabilities. Disable all unnecessary predefined accounts and, even though they are disabled, change the default passwords in case a future patch enables these accounts.
- Segregate reference design capabilities on their own subnetwork, separate from the production network, either physically or using virtual private networks and port-based authentication or similar mechanisms.
- Protect the various reference design subnetworks from each other and from the production network using security capabilities such as firewalls and intrusion detection devices that are configured per best practices.
- Configure firewalls to limit connections between the reference design network and the production network, except for connections needed to support required inter-network communications to specific IP address and port combinations in certain directions.

- Configure and verify firewall configurations to ensure that data transmission to and from reference design capabilities is limited to interactions that are needed. Restrict all permitted communications to specific protocols and IP address and port combinations in specific directions.
- Monitor the firewalls that separate the various reference design subnetworks from one another.
- NIST SP 1800-9C: *How-To Guides* contains the firewall configurations that show the rules implemented in each of the firewalls for the example implementation. These configurations are provided to enable the reader to reproduce the traffic filtering/blocking that was achieved in the implementation.
- Apply encryption or integrity-checking mechanisms to all information exchanged between reference design capabilities (i.e., to all user access, policy, and log information exchanged) so that tampering can be detected. Use only encryption and integrity mechanisms that conform to most recent industry best practices. Note that in the case of directory reads and writes, protected mode is defined as the use of Lightweight Directory Access Protocols (Request for Comments 2830).
- Strictly control physical access to both the reference design and the production network.
- Deploy a configuration management system to serve as a "monitor of monitors" to ensure that any changes made to the list of information are logged and reported to the monitoring system or to the analytics in the monitoring system and notifications are generated. Such a system could also monitor whether reference design monitoring capabilities, such as log integrity capabilities or the monitoring system itself, go offline or stop functioning, and generate alerts when these capabilities become unresponsive.
- Deploy a system that audits and analyzes directory content to create a description of who has access to what resources and validate that these access permissions correctly implement the enterprise's intended business process and access policies.

6.3.1.3 Policy Recommendations

- Define the access policies to enforce the principles of least privilege and separation of duties.
- Equip the monitoring capability with a complete a set of rules to take full advantage of the ability to identify anomalous situations that can signal a cyber event. Define enterprise-level work flows that include business and security rules to determine each user's access control authorizations and ensure that enterprise access control policy is enforced as completely and accurately as possible.
- Develop an attack model to help determine the type of events that should generate alerts.
- Grant only a very few users (e.g., human resource administrators) the authority to modify (initiate, change, or delete) employee access information. Require the approval of more than

915 one individual to update employee access information. Log all employee access information
916 modifications. Define work flows to enforce these requirements.

917 - Grant only a very few users (e.g., access rules administrators) the authority to modify (initiate,
918 change, or delete) access rules. Require the approval of more than one individual to update
919 access rules. Log all access rule modifications. Define work flows to enforce these requirements.

920 - Grant only a very few users (e.g., security analyst) the authority to modify (initiate, change, or
921 delete) the analytics that are applied to log information by the monitoring capability to
922 determine what constitutes an anomaly and generates an alert. Any changes made to the
923 analytics should, by policy, require the approval of more than one individual, and these changes
924 should themselves be logged, with the logs sent to a monitor-of-monitors system other than the
925 monitoring system and to all security analysts and other designated individuals. Define work
926 flows to enforce these requirements.

927 **Table 6-1 Capabilities for Managing and Securing the DI Reference Design**

928 This table describes only the product capabilities and CSF subcategory support used in the reference architecture. Many of the products have
929 significant additional security capabilities that are not listed here.

Capability	Specific Product	Function	CSF Subcategories
Subnetting	N/A	Technique of segmenting the network on which the reference design is deployed so that capabilities on one subnetwork are isolated from capabilities on other subnetworks. If an intruder gains access to one segment of the network, this technique limits the intruder's ability to monitor traffic on other segments of the network. For example, the enterprise's production network, on which user access information and decisions are conveyed, is separate from the reference design's monitoring and management subnetwork.	PR.DS-1: Data-at-rest is protected. PR.PT-4: Communications and control networks are protected.
Privileged Access Management	Active Directory	Manages privileged access to the OSs of all physical reference design capabilities. This is the single portal into which all users with administrator privileges must log in; it defines what systems these administrators are authorized to access based on their role and attributes. It also logs every login that is performed by users with administrator privileges, creating an audit trail of privileged	PR.AC-3: Remote access is managed. PR.AC-4: Access permissions are managed, incorporating the principles of least privilege and separation of duties. PR.PT-3: Access to systems and assets is controlled, incorporating the principle of least functionality.

Capability	Specific Product	Function	CSF Subcategories
		user access to the OSs of the physical systems that are hosting reference design capabilities.	DE.CM-3: Personnel activity is monitored to detect potential cybersecurity events.
Virtual Environment Privileged Access Management	Hyper-V VEEAM Active Directory	Manages privileged access to the virtual environment (including machines, switches, and host hardware) that host reference design capabilities. Hyper-V defines what VMs users are authorized to access based on the user's role. It logs activity that administrators perform on VMs, but it does not log operations that are performed on the OSs that are installed on those VMs. These logs create an audit trail of privileged user access to the virtual environment that is hosting the reference design capabilities.	PR.AC-3: Remote access is managed. PR.AC-4: Access permissions are managed, incorporating the principles of least privilege and separation of duties. PR.PT-3: Access to systems and assets is controlled, incorporating the principle of least functionality. DE.CM-3: Personnel activity is monitored to detect potential cybersecurity events.
Log Integrity	Tripwire Enterprise HPE ArcSight ESM	Forwards log information from each reference design capability to the monitoring capability. If an alternative product were used to instantiate this capability, it could add a time stamp and hash/integrity seal to each log file, thereby providing the file with integrity, but not confidentiality, protections. However, if the hash/integrity seal were to continue to be stored with the log file at the monitoring capability, it would provide a mechanism to	PR.DS-6: Integrity checking mechanisms are used to verify software, firmware, and information integrity. PR.PT-1: Audit/log records are determined, documented, implemented, and reviewed in accordance with policy. DE.AE-3: Event data is aggregated and correlated from multiple sources and sensors. PR.DS-2: Data-in-transit is protected.

Capability	Specific Product	Function	CSF Subcategories
		detect unauthorized modifications made to the log file while stored there.	

7 Functional Evaluation

A functional evaluation of the DI example implementation, as constructed in our laboratory, was conducted to verify that it meets its objective of demonstrating the ability to recover from DI attack. The evaluation verified that the example implementation could perform the following functions:

- recover from an identified ransomware attack
- recover from a data destruction event
- recover from a data manipulation event

Section 7.1 describes the format and components of the functional test cases. Each functional test case is designed to assess the capability of the example implementation to perform the functions listed above and detailed in Section 7.1.1.

7.1 Data Integrity Functional Test Plan

One aspect of our security evaluation involved assessing how well the reference design addresses the security characteristics it was intended to support. The CSF subcategories were used to provide structure to the security assessment by consulting the specific sections of each standard that are cited in reference to that subcategory. The cited sections provide validation points that the example solution is expected to exhibit. Using the CSF subcategories as a basis for organizing our analysis allowed us to systematically consider how well the reference design supports the intended security characteristics.

This plan includes the test cases necessary to conduct the functional evaluation of the DI example implementation, which is currently deployed in a lab at the NCCoE. The implementation tested is described in Section 5.

Each test case consists of multiple fields that collectively identify the goal of the test, the specifics required to implement the test, and how to assess the results of the test. Table 7-1 describes each field in the test case.

Table 7-1 Test Case Fields

Test Case Field	Description
Parent requirement	Identifies the top-level requirement or the series of top-level requirements leading to the testable requirement.
Testable requirement	Drives the definition of the remainder of the test case fields. Specifies the capability to be evaluated.
Associated security controls	Lists the NIST SP 800-53 rev 4 controls addressed by the test case.

Test Case Field	Description
Description	Describes the objective of the test case.
Associated test cases	In some instances, a test case may be based on the outcome of another test case(s). For example, analysis-based test cases produce a result that is verifiable through various means (e.g., log entries, reports, and alerts).
Preconditions	The starting state of the test case. Preconditions indicate various starting state items, such as a specific capability configuration required or specific protocol and content.
Procedure	The step-by-step actions required to implement the test case. A procedure may consist of a single sequence of steps or multiple sequences of steps (with delineation) to indicate variations in the test procedure.
Expected results	The expected results for each variation in the test procedure.
Actual results	The observed results.
Overall result	The overall result of the test as pass/fail. In some test case instances, the determination of the overall result may be more involved, such as determining pass/fail based on a percentage of errors identified.

7.1.1 Data Integrity Use Case Requirements

Table 7-2 identifies the DI functional evaluation requirements that are addressed in the test plan and associated test cases.

957 **Table 7-2 Data Integrity Functional Requirements**

Capability Requirement (CR) ID	Parent Requirement	Sub-requirement 1	Test Case
CR 1	The DI example implementation shall respond/recover from malware that encrypts files and displays notice demanding payment.		
CR 1.a		Produce notification of security event	Data Integrity -1
CR 1.b		Provide file integrity monitor	Data Integrity -1
CR 1.c		Revert to last known good	Data Integrity -1
CR 2	The DI example implementation shall recover when malware destroys data on user's machine.		
CR 2.a		Provide file integrity monitor	Data Integrity -2
CR 2.b		Revert to last known good	Data Integrity -2
CR 3	The DI example implementation shall recover when a user modifies a configuration file in violation of established baselines.		
CR 3.a		Provide file integrity monitor	Data Integrity -3 Data Integrity -6
CR 3.b		Revert to last known good	Data Integrity -3 Data Integrity -6
CR 3.c		Provide user activity auditing	Data Integrity -6

Capability Requirement (CR) ID	Parent Requirement	Sub-requirement 1	Test Case
CR 4	The DI example implementation shall recover when an administrator modifies a user's file.		
CR 4.a		Provide file integrity monitor	Data Integrity -4
CR-4.b		Provide user activity auditing	Data Integrity -4
CR 4.c		Revert to last known good	Data Integrity -4
CR-5	The DI example implementation shall recover when an administrator and/or script modifies data in a database.		
CR 5.a		Use database transaction auditing	Data Integrity -5
CR 5.b		Roll back to last known good	Data Integrity -5
CR-6	The DI example implementation shall recover when a user modifies a configuration file in violation of established baselines.		
CR 6.a		Provide file integrity monitor	Data Integrity -6
CR 6.b		Revert to last known good	Data Integrity -6
CR 6.c		Provide user activity auditing	Data Integrity -6

958

7.1.2 Test Case: Data Integrity-1

Table 7-3 Test Case ID: Data Integrity -1

Parent requirement	(CR 1) The **DI** example implementation shall respond/recover from malware that encrypts files and displays notice demanding payment.
Testable requirement	(CR 1.a) Logging, (CR 1.b) Corruption Testing, (CR 1.c) Backup Capability
Description	Show that the DI solution can recover from a DI attack that was initiated via ransomware.
Associated test cases	N/A
Associated CSF Subcategories	DE.DP-4, RS.CO-2, DE.EA-5, PR.DS-1, PR.DS-6, PR.PT-1
Preconditions	User downloaded and ran an executable from the internet that is ransomware. The user's files are then encrypted by the ransomware.
Procedure	1. Open the Tripwire Enterprise interface. 2. Click on the **Tasks** Section, enable the associated rule box, and click **Run.** 3. Open HPE ArcSight ESM. 4. Under Events, select **Active Channels**, then select **Audit Events.** 5. Find the Tripwire Enterprise event logs associated with the event. Select **Fields** in the **Customize** dropdown and enable the following fields: a. **End Time** b. **Attacker Address** c. **File Name** d. **Device Action** e. **Source User Name** f. **Device Custom String6** 6. Open IBM Spectrum Protect. 7. Click on **Restore.** 8. Select missing files and click **Restore to original location.**
Expected Results (pass)	Event identified (CR 1.a) Details of the event are understood and moment of last known good is identified.

	Provide file Integrity monitor (CR 1.b). Modified files are correctly identified. Recovery complete (CR 1.c). System was restored to pre-DI event version.
Actual Results	Details of the event were understood and the moment of last known good was identified for the file in question. All the files affected within that timeframe were correctly identified, and a full and successful restore was executed.
Overall Result	Pass. All metrics of success were met to satisfaction.

7.1.3 Test Case Data Integrity-2

Table 7-4 Test Case ID: Data Integrity -2

Parent requirement	(CR 2) The **DI** example implementation shall recover when malware destroys data on user's machine.
Testable requirement	(CR 2.a) Corruption Testing, (CR 2.b) Backup Capability
Description	Show that the DI solution can recover from a DI attack that destroys data via a malware attack.
Associated test cases	N/A
Associated CSF Subcategories	PR.DS-1, PR.IP-4, PR-DS-6, PR.PT1
Preconditions	User downloads a malicious executable that modifies critical data.
Procedure	1. Open the Tripwire Enterprise interface. 2. Click on the **Tasks** Section, enable the associated rule box, and click **Run.** 3. Open HPE ArcSight ESM. 4. Under Events, select **Active Channels**, then select **Audit Events.** 5. Find the Tripwire event logs associated with the event. Select **Fields** in the **Customize** dropdown and enable the following fields: a. **End Time** b. **Attacker Address** c. **File Name** d. **Device Action** e. **Source User Name** f. **Device Custom String** 6. Open IBM Spectrum Protect. 7. Click on **Restore.** 8. Select missing files and click **Restore to original location.**
Expected Results (pass)	Provide file integrity monitor (CR 2.a). Modified files are correctly identified. Recovery complete (CR 2.b).

	System was restored to pre-DI event version.
Actual Results	Details of the event were understood and the moment of last known good was identified for the file in question. All the files affected within that timeframe were correctly identified, and a full and successful restore was executed.
Overall Result	Pass. All metrics of success were met to satisfaction.

7.1.4 Test Case Data Integrity-3

Table 7-5 Test Case ID: Data Integrity -3

Parent requirement	(CR 3) The **DI** example implementation shall recover when a user modifies a configuration file in violation of established baselines.
Testable requirement	(CR 3.a) Corruption Testing, (CR 3.b) Backup Capability
Description	Show that the DI solution can recover from a DI event that modifies system configurations.
Associated test cases	N/A
Associated CSF Subcategories	PR.DS-1, PR.DS-6, PR.PT-1, DE.CM-3, DE.AE-1, DE.CM-1
Preconditions	Run a script that would simulate the effects of a configuration modification event.
Procedure	1. Open HP ArcSight ESM. 2. Under **Events**, select **Event Search.** 3. Use the search bar to search for the keyword “created” to find associated event logs for account creation. 4. After determining the point in time of a malicious event, restart the Active Directory server, holding down the **F2** and **F8** keys while restarting to enter the **Advanced Boot Options** menu. 5. Select **Directory Services Repair Mode.** 6. Log in as the machine administrator. 7. Open a command prompt. 8. View visible backup versions with the following command: ▪ **`wbadmin get versions`** 9. Restore to a selected backup target with the following command. Note that the selected date should reflect the last known good backup: ▪ **`wbadmin start systemstaterecovery -version:<Version Number> -backupTarget:<Backup Location>`** ▪ Replace **`<Version Number>`** with the desired version’s version identifier, and **`<Backup Location>`** with the version’s corresponding backup location.

	10. Provide a username (with domain if applicable) and password for a privileged user to the backup location. 11. Acknowledge the remaining prompts and wait for the backup to complete. The system will automatically restart.
Expected Results (pass)	Provide file integrity monitor (CR 3.a). Modified files are correctly identified. Recovery complete (CR 3.b). Modified files are restored to their original state.
Actual Results	The fake accounts were successfully identified and deleted. The remaining accounts were restored to their original states at the time of the backup.
Overall Result	Pass. All metrics of success were met to satisfaction.

7.1.5 Test Case Data Integrity-4

Table 7-6 Test Case ID: Data Integrity -4

Parent requirement	(CR 4) The **DI** example implementation shall recover when an administrator modifies a user's file.
Testable requirement	(CR 4.a) Corruption Testing, (CR 4.b) Logging, (CR 4.c) Backup Capability
Description	Show that the DI solution can recover from when an administrator modifies a user's file.
Associated test cases	N/A
Associated CSF Subcategories	DE.AE-1, DE.AE-3, DE.AE-5
Preconditions	Two VMs on Microsoft Hyper-V have been backed up. Administrator accidentally runs a command that deletes a critical VM. `Remove-VM -Name "<VMName>" -Force`
Procedure	1. Open HP ArcSight ESM. 2. Under **Events**, select **Event Search.** 3. Use the search bar to search for the deleted VM's name and then find the associated event log. 4. Locate previous logins from that machine by searching for the VM host machine's domain and name in the search bar. Look for logins before the time of the deletion incident, without an associated logout before the event. User logins (as opposed to automated ones that occur constantly in the machine) will have a non-null value for the **Source Address** field, typically 127.0.0.1. 5. Open the VEEAM console. 6. Navigate to the **Backups** menu. 7. Right-click on deleted VM and click **Restore**, and then **Entire VM.** 8. When prompted, search for the deleted VM's name and select it for restoration. 9. When prompted, enter reason for VM restoration.
Expected Results (pass)	Provide file integrity monitor (CR 4.a). Missing files are correctly identified.

	Provide user activity auditing (CR 4.b). User who initiated deletion is correctly identified. Revert to last known good (CR 4.c). VM is fully restored to original functionality.
Actual Results	The VEEAM system functioned as expected. Deleted VM is restored to its original functionality. Any user logged in during the deletion event was identified.
Overall Result	Pass (partial). The file integrity monitoring and reversion to last known good requirements were met. User activity was audited, but it is not possible to determine which user caused the deletion event if multiple users were logged in to the machine at the time of the event.

7.1.6 Test Case Data Integrity-5

Table 7-7 Test Case ID: Data Integrity -5

Parent requirement	(CR 5) The **DI** example implementation shall recover when an administrator and/or script modifies data in a database.
Testable requirement	(CR 5.a) Logging, (CR 5.b) Backup Storage
Description	Show that the DI solution can recover when data in a database has been altered in error by an administrator or script.
Associated test cases	N/A
Associated CSF Subcategories	DE.AE-3, DE.AE-5
Preconditions	Run a script that would simulate the effects of an administrator or script modification within a database.
Procedure	1. Open HP ArcSight ESM. 2. Under **Events**, select **Event Search.** 3. Use the search bar to search for the affected database and then find the associated event log. Use the field **cs1** to find the affected table name and **cs2** to find the undesired database transaction query string. Modify time parameters for the search to narrow the desired transaction. 4. Use the **duser** field of the event to find the name of the user who executed the transaction event. 5. Determine the number of transactions that occurred and then use a transactional rollback tool to restore the database to the last known good state.
Expected Results (pass)	Use database transaction auditing (CR 5.a). Bad database transaction is correctly identified. Roll back to last known good (CR 5.b). Database is restored to full functionality.

Actual Results	The database data was successfully restored to its last known good state. The user responsible for the event was identified and the time of the event was determined.
Overall Result	Pass. All metrics of success were met to satisfaction.

7.1.7 Test Case Data Integrity-6

Table 7-8 Test Case ID: Data Integrity -6

Parent requirement	(CR 6) The **DI** example implementation shall recover when a user modifies a configuration file in violation of established baselines.
Testable requirement	(CR 6.a) Corruption Testing, (CR 6.b) Backup Capability (CR 6.c). Provide user activity auditing.
Description	Show that the DI solution can recover when the database schema has been altered in error by an administrator or script.
Associated test cases	N/A
Associated CSF Subcategories	PR.DS-1, PR.DS-6, PR.PT-1, DE.CM-3, DE.AE-1, DE.CM-1
Preconditions	Run a script that would simulate the effects of an administrator or script modifying the database schema.
Procedure	1. Open the Tripwire Enterprise interface. 2. Click on the **Tasks** Section, enable the associated rule box, and click **Run.** 3. Open HP ArcSight ESM. 4. Under Events, select **Active Channels**, then select **Audit Events.** 5. Find the Tripwire event logs associated with the event. Select **Fields** in the **Customize** dropdown and enable the following fields: a. **End Time** b. **Attacker Address** c. **File Name** d. **Device Action** e. **Source User Name** f. **Device Custom String6** 6. Open SQL Server Management Studio and locate the affected database(s).

	7. Right-click on the database name and select **Tasks > Restore > Database...** 8. Verify that the **Restore To:** location is a backup from before the time of the incident. 9. Under **Options,** select **Overwrite the existing database (WITH REPLACE)** 10. Click **OK** and wait for the restoration to complete.
Expected Results (pass)	Provide file integrity monitor (CR 6.a). Modified table is correctly identified. Revert to last known good (CR 6.b). Database fully restored to previous functionality. Provide user activity auditing (CR 6.c). User who initiated the modification is correctly identified.
Actual Results	The database schema was successfully restored to its last known good state. The user responsible for the event was identified and the time of the event was determined.
Overall Result	Pass. All metrics of success were met to satisfaction.

971

8 Future Build Considerations

The NCCoE is considering additional DI projects that map to the Cybersecurity Framework Core Functions of Identify, Protect, Detect and Respond. This reference design focuses largely on the Recover aspect of the CSF. The functions of the CSF lead into each other and act as a cycle. Identifying vulnerabilities leads to protection against them. Protecting against vulnerabilities allows enterprises to detect cybersecurity events. Detection of events gives enterprises the information needed to respond and recover from these events as well as reshape their policy to identify and protect against events in the future. Though this project deals primarily with an organization's capabilities to recover from DI events, future NCCoE projects may look at capabilities for meeting the requirements of the other functions in the CSF.

This project does not include instructions for automated full system recovery. If malicious software manages to affect critical system files, recovery becomes more difficult. The backup software used is client-based, so the system must be able to run the client to restore, which may not be possible in some instances. Solutions exist to help automate the process to fully restore a failed system and integrate with existing backup solutions. A future build might include the use of a product to address these types of attacks.

This project uses built-in database capabilities to achieve transactional rollbacks as well as database metadata restoration. The restoration process is granular and uses built-in mechanisms; however, automating the process is more difficult. Products exist that use the built-in restoration mechanisms and implement their own database backup functionality. These products add varying degrees of latency to database transactions, depending on the mechanisms used and the granularity of recovery the organization desires.

Appendix A List of Acronyms

COI	Community of Interest
CR	Capability Requirement
CSF	Cybersecurity Framework
DI	Data Integrity
ESM	Enterprise Security Manager
HPE	Hewlett Packard Enterprise
IEC/ISO	International Electrotechnical Commission/International Organization for Standardization
IP	Internet Protocol
IT	Information Technology
MS SQL	Microsoft Structured Query Language
NCCoE	National Cybersecurity Center of Excellence
NIST	National Institute of Standards and Technology
OS	Operating System
SP	Special Publication
VM	Virtual Machine
WORM	Write Once Read Many

Appendix B References

[1] A. Sedgewick, *Framework for Improving Critical Infrastructure Cybersecurity*, Version 1.0, National Institute of Standards and Technology, Gaithersburg, Maryland, February 2014, 41pp. https://www.nist.gov/sites/default/files/documents/cyberframework/cybersecurity-framework-021214.pdf [accessed 7/10/17]

[2] L. Kauffman and B. Abe, *Executive Technical Workshop on Improving Cybersecurity and Consumer Privacy*, NISTIR 8050, National Institute of Standard and Technology, Gaithersburg, Maryland, April 2015, 15pp. https://nccoe.nist.gov/sites/default/files/library/nistir-8050-draft.pdf [accessed 7/10/17]

[3] G. Stoneburner *et al.*, *Guide for Conducting Risk Assessments*, NIST Special Publication (SP), 800-30 Revision 1, National Institute of Standards and Technology, Gaithersburg, Maryland, September 2012, 95pp. http://dx.doi.org/10.6028/NIST.SP.800-30r1

[4] R. Ross *et al.*, *Guide for Applying the Risk Management Framework to Federal Information Systems*, NIST Special Publication (SP) 800-37, National Institute of Standards and Technology, Gaithersburg, Maryland, February 2010, 101pp. http://dx.doi.org/10.6028/NIST.SP.800-37r1

[5] R. Ross *et al.*, *Managing Information Security Risk*, NIST Special Publication (SP) 800-39, National Institute of Standards and Technology, Gaithersburg, Maryland, March 2011, 87pp. http://dx.doi.org/10.6028/NIST.SP.800-39

[6] M. Souppaya *et al.*, *Guide to Enterprise Patch Management Technologies*, NIST Special Publication (SP) 800-40 Revision 3, National Institute of Standards and Technology, Gaithersburg, Maryland, July 2013, 25pp. http://dx.doi.org/10.6028/NIST.SP.800-40r3

[7] R. Ross *et al.*, *Security and Privacy Controls for Federal Information Systems and Organizations*, NIST Special Publication (SP) 800-53 Revision 4, National Institute of Standards and Technology, Gaithersburg, Maryland, April 2013, 461pp. https://doi.org/10.6028/NIST.SP.800-53r4

[8] U.S. Department of Commerce. *Security Requirements for Cryptographic Modules*, Federal Information Processing Standards (FIPS) Publication 140-2, May 2001, 69pp. http://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.140-2.pdf [accessed 8/4/17].

[9] K. Kent *et al.*, *Guide to Integrating Forensic Techniques into Incident Response*, NIST Special Publication (SP) 800-86, National Institute of Standards and Technology, Gaithersburg, Maryland, August 2006, 121pp. http://dx.doi.org/10.6028/NIST.SP.800-86

[10] K. Kent and M. Souppaya, *Guide to Computer Security Log Management*, NIST Special Publication (SP) 800-92, National Institute of Standards and Technology, Gaithersburg, Maryland, September 2006, 72pp. http://dx.doi.org/10.6028/NIST.SP.800-92

[11] P. Bowen *et al.*, *Information Security Handbook: A Guide for Managers*, NIST Special Publication (SP) 800-100, National Institute of Standards and Technology, Gaithersburg, Maryland, October 2006, 178pp. http://dx.doi.org/10.6028/NIST.SP.800-100

[12] M. Swanson *et al*., *Contingency Planning Guide for Federal Information Systems*, NIST Special Publication (SP) 800-34 Revision 1, National Institute of Standards and Technology, Gaithersburg, Maryland, May 2010, 148pp. http://dx.doi.org/10.6028/NIST.SP.800-34r1

[13] Office of Management and Budget (OMB), *Management of Federal Information Resources,* OMB Circular No. A-130, November 2000. https://www.whitehouse.gov/omb/circulars_a130_a130trans4 [accessed 8/4/17].

[14] P. Cichonski *et al.*, *Computer Security Incident Handling Guide*, NIST Special Publication (SP) 800-61 Revision 2, National Institute of Standards and Technology, Gaithersburg, Maryland, August 2012, 79pp. http://dx.doi.org/10.6028/NIST.SP.800-61r2

[15] M. Souppaya and K. Scarfone, *Guide to Malware Incident Prevention and Handling for Desktops and Laptops*, NIST Special Publication (SP) 800-83 Revision 1, National Institute of Standards and Technology, Gaithersburg, Maryland, July 2013, 46pp. http://dx.doi.org/10.6028/NIST.SP.800-83r1

[16] C. Johnson *et al.*, *Guide to Cyber Threat Information Sharing*, NIST Special Publication (SP) 800-150, National Institute of Standards and Technology, Gaithersburg, Maryland, October 2016, 42pp. http://dx.doi.org/10.6028/NIST.SP.800-150

[17] M. Bartock *et al.*, *Guide for Cybersecurity Event Recovery*, NIST Special Publication (SP) 800-184, National Institute of Standards and Technology, Gaithersburg, Maryland, December 2016, 52pp. http://dx.doi.org/10.6028/NIST.SP.800-184