

Technical Validation

Protecting Critical Data from Cyber Threats Such As Ransomware with a Comprehensive Digital Vault Solution

Dell EMC PowerProtect Cyber Recovery with CyberSense

By Vinny Choinski, Senior Validation Analyst; and Christophe Bertrand, Senior Analyst
October 2020

This ESG Technical Validation was commissioned by Dell Technologies and is distributed under license from ESG.



Contents

Introduction..... 3

 Background..... 3

 Dell EMC PowerProtect Cyber Recovery Solution..... 4

ESG Technical Validation 5

 PowerProtect Cyber Recovery Vault..... 5

 Cyber Recovery Air Gap 7

 CyberSense 8

The Bigger Truth 10

ESG Technical Validations

The goal of ESG Technical Validations is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Technical Validations are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team’s expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.

Introduction


This ESG Technical Validation documents hands-on validation of the Dell EMC PowerProtect Cyber Recovery solution with CyberSense from Dell Technologies. ESG performed analysis and auditing of all major elements of the Cyber Recovery solution, including the Cyber Recovery vault and its multiple layers of security and protection for critical data, the operational air gap, and the CyberSense analytics engine.

Background

Regardless of the industry or size of the organization, cyber-attacks continually expose business and governments to compromised data, lost revenue due to downtime, reputational damage, and costly regulatory fines.

ESG research confirms that most organizations will increase cybersecurity spending in 2020, driven by the desire to protect business processes and counteract dangerous threats. While most are likely to invest in AI/ML-based analytics, data security, network security, and application security, CISOs will spread budget dollars around in many areas (see Figure 1). The data indicates that many organizations are in the process of reengineering their entire cybersecurity infrastructure in an attempt to improve efficacy, streamline security operations, and support new technology-driven business processes.¹

Figure 1. Top 7 2020 Cybersecurity Investment Areas



Source: Enterprise Strategy Group

Ransomware and bad actors have learned that they have a better chance of achieving their goal—getting paid a ransom or destroying data—if no backup is available for recovery. This has placed backup systems and data under direct attack. Unfortunately, there are a number of vulnerable points because backups were designed for accessibility, not necessarily for security. Data and application backup systems become targets for compromise because access to these systems is trusted and can therefore be more easily and rapidly infiltrated. Plus, backup systems often contain data catalogs, and encrypting or deleting this catalog data can make recovery from a backup take much longer—or even prevent the backup from being used for recovery.


¹ Source: ESG Brief, [2020 Cybersecurity Spending Trends](#), March 2020.

Dell EMC PowerProtect Cyber Recovery Solution

Cyber Recovery provides a comprehensive solution to combat ransomware and destructive cyber-attacks. It is a mature solution with five years in the market and hundreds of customers. The Cyber Recovery software automates synchronization of data between production systems and the vault, creating air-gapped, immutable copies with locked retention policies. If a cyber-attack occurs, users can quickly identify a clean copy of data, recover critical systems, and get the business operational again.

PowerProtect Cyber Recovery relies on the concept of a data “vault,” which is often physically isolated—within a locked cage or room—and always logically isolated via an operational air gap. The vault is not an extra data center—it is usually located at the production or corporate data center or sometimes with a third-party solution provider. The vault components are never accessible from production. Access to the vault is provided through an “operational air gap,” as shown in Figure 2.

Figure 2. Overview of Dell EMC PowerProtect Cyber Recovery Solution



Source: Enterprise Strategy Group

Key elements of the Cyber Recovery solution include:

- Only the PowerProtect Cyber Recovery solution has an automated, orchestrated logical air gap to protect data. Some other solutions claim an “air gap” because the data can be separate from the production network (for example, a copy of the data stored in the cloud); however, with this approach the data could still be accessible to bad actors and an offsite copy alone does not provide a complete air gap.
- The Cyber Recovery vault itself is physically and logically isolated. It cannot be opened or controlled from the production side nor accessed unless the person is physically in the vault.
- The PowerProtect DD Compliance Retention Lock capability complies with the 17a-4(f)(ii) standard. It cannot be turned off without a separate security officer password, and then only for newly stored data. This capability has been further hardened to protect against even more advanced attacks such as NTP manipulation.
- Cyber Recovery can protect and enable the recovery of any vendor’s backup sets that can write to the PowerProtect appliance.
- In August 2020, PowerProtect Cyber Recovery became the first, and at the time that this paper was published, the only solution to receive endorsement for meeting all of the data vaulting requirements of the [Sheltered Harbor](#) standard.

As part of PowerProtect Data Manager and for users running Dell EMC NetWorker, PowerProtect Cyber Recovery enables automated recovery from the vault. Dell and its ecosystem partners provide a comprehensive methodology for protecting data as well as performing damage assessments and forensics to recover systems or remediate and remove the offending malware.


ESG Technical Validation

ESG performed validation and testing of the PowerProtect Cyber Recovery solution using the Dell Proof of Concept Engineering Lab. Testing was designed to validate all major elements of the Cyber Recovery solution, including the Cyber Recovery vault, the operational air gap, and the CyberSense analytics engine.

PowerProtect Cyber Recovery Vault

The PowerProtect Cyber Recovery solution relies on the deployment and operation of a Cyber Recovery “vault” as its foundational element. The vault offers multiple layers of protection to provide resilience against cyber-attacks, even from an insider threat. It moves critical data away from the attack surface, physically isolating it within a protected part of the data center and requires separate security credentials and multifactor authentication for access. Additional safeguards include an automated operational air gap to provide network isolation, plus the elimination of production-accessible management interfaces that can be compromised. As illustrated in Figure 3, Cyber Recovery is managed through an intuitive interface, accessible only from the vault, complete with comprehensive solution status dashboards. A CLI interface and APIs are also available for management. Cyber Recovery software deployed within the security of the vault environment automates the synchronization of data between production systems and the vault and then creates immutable copies with locked retention policies. If a cyber-attack occurs, a clean data copy can quickly be identified, and business-critical systems recovered.

Figure 3. PowerProtect Cyber Recovery Management Interface Dashboards



Source: Enterprise Strategy Group

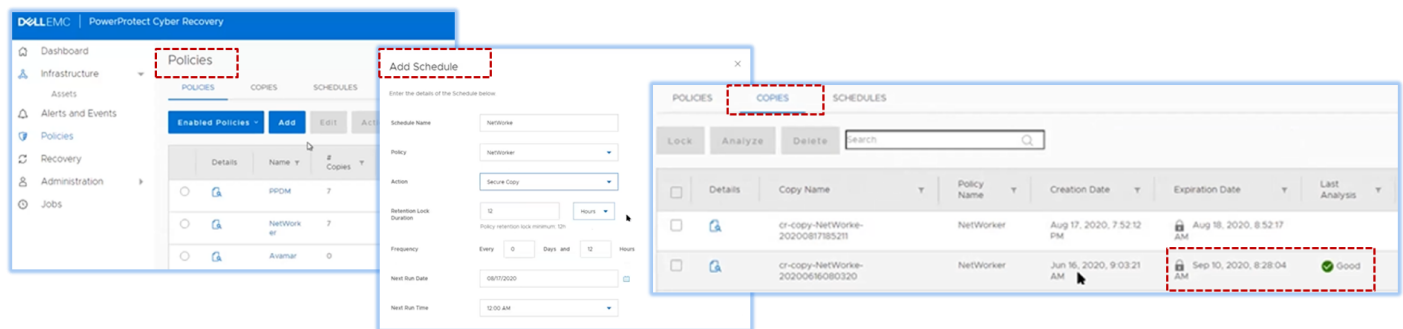
ESG exercised and validated the standard operational data flows and application performance of systems related to the Cyber Recovery solution, as well as those within the vault itself. Users would back up production data and applications in a solution such as the PowerProtect DD and NetWorker environment shown in the reference architecture, though Cyber Recovery can interface with many different vendors’ backup applications.

The air gap feature will be discussed in greater detail in the next section of this report, but essentially, a specific, preconfigured port(s) in the vault firewall is designated for data ingest and then the Cyber Recovery software enables the port at short, policy-driven intervals. Data is “pulled” by the vault systems from the backup solution, through the air gap, and into the vault’s PowerProtect DD target. Once the data copy has been ingested, the designated port is disabled, effectively re-sealing the vault. By utilizing an SEC 17a-4(f)(ii) compliant retention lock and WORM immutability, and requiring separate security officer and multifactor access credentials, the Cyber Recovery solution design makes it

essentially impossible for the vault to be hacked. ESG verified that Cyber Recovery incorporates multiple security layers to protect against intruders, including insiders. The vault cannot be opened or controlled from the production side. Vault deployments utilize both physical and logical isolation—properly deployed vault systems can't be accessed unless the person is physically in the vault. Data retention locks can't be turned off without a separate security officer password.

As illustrated in Figure 3, ESG used the management UI to easily create automated Cyber Recovery policies. Once the PowerProtect DD target ingests the data and the air gap has been closed, Cyber Recovery makes the copy and related catalog immutable by placing a policy-driven retention lock on the files to protect them from accidental or intentional deletion. As noted, vault retention is configurable, but most users keep about a month's worth of copies.

Figure 4. PowerProtect Cyber Recovery Policy Management



Source: Enterprise Strategy Group

After the data is secured, it can be scanned by the integrated CyberSense analytics engine to uncover anomalies. ESG validated that CyberSense employs full file content analytics and machine learning (ML) trained on numerous attack vectors to identify possible data compromise. It operates securely in the vault, where it evaluates the data stored there, taking over 100 observations per file each day. This data is then evaluated within the vault by an ML/AI engine to generate a verdict on the validity of the data. As with the air gap feature, the CyberSense analytics capabilities and functionality will be discussed in greater detail later in this report.

Recovering data from the vault in the event of a cyber-attack or simply for recovery testing procedures is critical. ESG verified that the Cyber Recovery solution provides a number of ways recovery can be performed. The solution includes in-vault intelligence tools to accelerate recovery of “clean” copies. CyberSense renders a verdict on each data set to determine whether it is “OK” or “Suspicious.” Data labeled “Suspicious” includes information about content that is at risk, so that the rest of the “clean” data can be used for recovery. This enables the fastest and surest possible recovery, in addition to potentially helping to track the source of the attack. Dell claims that no competitive solution provides specifics, during a first pass analysis, on portions of the data set that may have been impacted; and several competitors’ solutions rely on information being sent to or extracted from a public cloud SaaS plane that is unlikely to be available from the production environment, which is normally shut down post-attack.

Why This Matters

Many organizations have accepted the fact that they will experience a successful cyber-attack at some point in the not-too-distant future. These organizations have shifted their thinking from prevention to resilience—they want to recover business-critical systems as quickly and efficiently as possible after a cyber incident.


The PowerProtect Cyber Recovery solution makes a “vault” approach to cyber resilience practical. ESG validated that deployment and operation of the vault solution and its various components is relatively straightforward, especially with the help of Dell EMC Advisory Services.

Cyber Recovery Air Gap

The concept of an “air gap” is crucial to the “vault” approach utilized by the PowerProtect Cyber Recovery solution. Creating a successful air gap that truly isolates the vault environment from external systems and connectivity requires careful engineering and innovative solution design. Dell calls their implementation an “operational air gap” to differentiate from the older “air gap” definition which identifies a network that is never externally connected. Cyber Recovery’s air gap is also “operational” because ingestion of data and management of the process is automated and policy-driven and requires no manual intervention.

As illustrated in Figure 5 and verified by ESG, the vault and its systems and data are physically and logically isolated except during the actual data ingest process. From the security of the vault, the PowerProtect Cyber Recovery software initiates a scheduled data transfer by designating, configuring, and briefly enabling a single network port on an interface dedicated for that purpose. In the Dell Proof-of-Concept Engineering Lab, data was transferred between a PowerProtect DD 9900 system used as a backup device on the production side, across a briefly enabled network connection, to a PowerProtect DD 9900 target system in the vault. Data was “pulled” from the backup system on the production side, so it did not participate in management or operation of the network connection. As soon as the data transfer was complete, the network connection was not only closed, but the interface itself was disabled, thus reestablishing complete vault isolation via an operational air gap. Air gap management was entirely controlled by the Cyber Recovery software from within the vault. And even when the air gap was briefly unlocked, access is extremely limited—only the replication destination is ever accessible, to accept the updated data, and vault components, including prior data copies, are never accessible.

Figure 5. PowerProtect Cyber Recovery Operational Air Gap



Source: Enterprise Strategy Group

Why This Matters

The Cyber Recovery vault is a powerful solution that enhances organizational cyber resilience. But the vault is only as effective as the air gap technology and approach used to establish and maintain isolation. ESG confirmed that the operational air gap solution used by Cyber Recovery is state-of-the-art, performs its intended function, and provides a secure solution for protecting critical data from cyber-attacks, ransomware, malware, and other threats that ESG has seen in the market.

For all the organizations and IT decision makers evaluating various cyber resilience solutions, having confidence in the ability to isolate data from ongoing cyber intrusion is a key ingredient in a successful cyber resilience solution.

CyberSense

ESG validated the basic configuration and operation of the CyberSense analytics engine and functionality within the PowerProtect Cyber Recovery solution environment as implemented in the Dell Proof-of-Concept Engineering Lab. CyberSense is integrated into the overall Cyber Recovery solution and adds an intelligent layer of protection to help find data corruption if a cyber-attack penetrates the data center production or backup environments. The CyberSense analytics engine provides full content indexing and uses machine learning to analyze over 100 content-based statistics to detect signs of corruption due to cyber intrusions. Dell claims that CyberSense finds corruption with up to 99.5% confidence, helping users identify threats and diagnose attack vectors while protecting business-critical content, all within the security of the vault. Figure 6 illustrates the basic steps involved in the CyberSense analytics and reporting process.

Figure 6. CyberSense Analytics and Reporting Process



Source: Enterprise Strategy Group

- Data is scanned in the format in which it was stored in the vault—e.g., backups remain in backup format.
- Analytics take over 100 observations per file, including measurements such as entropy and similarity.
- These observations are collected and evaluated by a machine learning tool that has identified patterns indicating data has been corrupted. Because it is looking for patterns, not signatures, the analysis is more effective and does not need to be updated as frequently. The process is repeated each time a new data set is brought into the vault.
- Reporting and analytics tools can help to quickly identify corrupt data and the source of an attack, enabling faster and more assured recovery.

Once data is replicated to the Cyber Recovery vault and a retention lock is applied, CyberSense scans the backup data, creating point-in-time observations of files and databases. Running analytics on the data in the vault is an important component to both maintain integrity of this data and enable a speedy recovery after an attack. Analytics help to determine whether a data set is valid and usable for recovery or has somehow been improperly altered or corrupted so that it's "Suspicious" and potentially unusable.

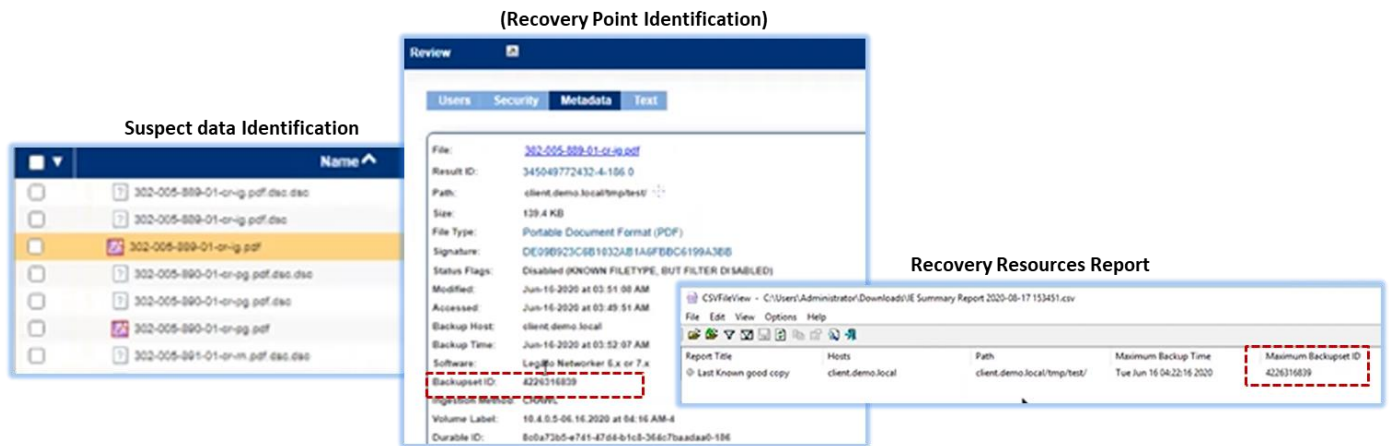
CyberSense detects mass deletions, encryption, and other types of changes in files and databases that result from common cyber-attacks. This scan occurs directly on the data within the backup image without the need for the original backup software. The analytics generated include file type mismatch, corruption, known ransomware extensions, deletions, entropy, similarity, and more. The analytics results are then used by machine learning algorithms to make a deterministic decision on data corruption that is indicative of a cyber-attack. The machine learning algorithms have been trained by all the latest trojans and ransomware and can be updated as new attack vectors are discovered. CyberSense analytics are particularly powerful because (a) they can read through the backup format so there is no need to restore data, which is difficult to automate and can expose the vault to potential risk; and (b) they evaluate the full contents of the file, not just its metadata, to deliver superior insights.

CyberSense goes beyond metadata-only solutions because it is based on analyzing the full content of the vaulted files, which provides a high degree of confidence in detecting data corruption. It audits files and databases for attacks that include content-only based corruption of the file structure or partial encryption inside a document or page of a database. These attacks cannot be found using analytics that do not scan inside the file to compare how it is changing over time. Without full content-based analytics, the number of false negatives will be significant, providing a false sense of confidence in data integrity and security.

Ongoing observations of the data allow CyberSense to track how the contents of files change over time. CyberSense generates analytics from a comprehensive range of data types, including core infrastructure such as DNS, LDAP, and Active Directory; unstructured data files such as documents, contracts, and agreements; intellectual property; and databases such as Oracle, DB2, SQL, Epic Caché, and others.

If CyberSense detects signs of corruption, an alert is generated that includes the attack vector and listing of files affected. A critical alert is displayed in the Cyber Recovery dashboard, then CyberSense post-attack forensic reports are available to quickly diagnose and recover from the ransomware attack. As illustrated in Figure 7, with CyberSense, organizations can proactively audit their files and databases to determine when an attack began, then quickly recover with the last good version of the data, perhaps even before there is any interruption to the business.

Figure 7. CyberSense Remediation



Source: Enterprise Strategy Group

Why This Matters

It's one thing to isolate copies of crucial data from the rest of the data center. It's a very different thing to have confidence that those copies can be used to safely recover business operations. The confidence that particular data sets really can be used to recover business-critical systems is what matters most to the organizations deploying cyber resilience solutions such as PowerProtect Cyber Recovery.

CyberSense analytics complete the Cyber Recovery process by offering confidence to users that the data contained in the Cyber Recovery vault is "clean" and usable for actual recovery purposes. The full content and machine learning-driven data analysis are state-of-the-art for this type of operation, so organizations that depend on PowerProtect Cyber Recovery feel confident that they are doing everything they can to increase their cyber resilience and lower the risk and cost of cyber-attacks and intrusions.

The Bigger Truth

Organizations are moving away from a focus on cybersecurity and toward a more comprehensive approach known as cyber resilience. Rather than hope that they can defeat cyber threats 100% of the time, they are instead assuming that a cyber-attack will eventually succeed and the organization will need to recover from the attack and resume operations as quickly and cost-efficiently as possible.

In order to successfully recover business-critical data and systems after a cyber incident, unaffected or “clean” copies of applications and data sets must be available. The challenge becomes how an organization ensures that clean data and unaffected applications exist. For many organizations, the answer is to physically and logically isolate data copies away from all normal production and backup systems, then scan and analyze this data inside the secure “vault” environment to detect any corruption and identify clean copies that can be used with confidence for operational recovery.

The PowerProtect Cyber Recovery solution from Dell Technologies addresses an important requirement within the rapidly expanding cyber resilience marketplace by providing a complete digital vault, data analysis, and cyber recovery solution. Cyber Recovery integrates key elements of an overall solution, including an automated, state-of-the-art “air gap” protection, plus a proven, AI-driven analytics engine that operates from the security of the vault environment.

Cyber resilience strategies often employ terms such as “trust” and “confidence” to describe the levels of effectiveness provided by one solution or another. ESG hands-on evaluation confirmed that organizations deploying the PowerProtect Cyber Recovery solution can indeed tell shareholders that they have a high degree of confidence that systems can successfully be recovered after a cyber-attack and the business will be operational and profitable again with the least amount of disruption.

Confidence is crucial in the marketplace. This is what PowerProtect Cyber Recovery offers—proven and modern data protection that improves cyber resilience. It’s exactly what many organizations are looking for.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.



Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.

© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.

