



[www.esti.ca](http://www.esti.ca)

Head Office  
302 2<sup>nd</sup> Ave. N.  
Saskatoon, SK  
S7K 2B9  
T: (306) 242-2436  
F: (306) 242-7329

1939 Scarth Street  
Unit 201  
Regina, SK  
S4P 2H1  
T: (306) 546-0716

Winnipeg, MB  
T: (204) 927-1932

10180 101 Street NW  
Suite 3400  
Edmonton, AB  
T5K 3S4  
T: (780) 628-6119

Calgary, AB  
T: (403) 770-9141

Vancouver, BC  
T: (778) 785-0270

*Your Problems.*  
*Our Solutions.*

# ESTI Consulting Services Technical Whitepaper

## Dell PowerProtect Cyber Recovery Vault on Microsoft Azure

March 25, 2023 – Version 1.0

Created by:  
Earl Gosick – Partner, Storage Specialist  
ESTI Consulting Services



## Table of Contents

<b>1. Introduction</b>	<b>3</b>
<b>2. Dell EMC PowerProtect Cyber Recovery Vault</b>	<b>3</b>
High Level Architecture	4
Key Components	5
Optional Components to enhance security	5
Key features of the Cyber Recovery software	6
<b>3. Why is Dell EMC PowerProtect Cyber Recovery Vault a more effective solution?</b>	<b>6</b>
Cyber Recovery Air Gap	6
Full content analytics of data within the Vault	7
Ability to Recover following an attack	8
Deployment options for a Cyber Recovery Vault	8
<b>4. Dell PowerProtect Cyber Recovery Vault deployed on Microsoft Azure</b>	<b>9</b>
PowerProtect Cyber Recovery for Azure	9
PowerProtect Cyber Recovery and DDVE	10
Architecture overview	11
Assigned resources in the Cyber Recovery Azure virtual network (VNet)	12
Cyber Recovery on Azure limitations and unsupported features	13
<b>5. ESTI Consulting Services</b>	<b>13</b>
<b>6. Summary</b>	<b>14</b>
<b>7. References and Additional Documentation</b>	<b>15</b>

## 1. Introduction

Modern organizations increasingly rely on data, and how they can leverage that data, to increase success. Machine Learning and Artificial intelligence rely on data to effectively provide insight supporting business goals. This trend is driving an increase in the value of all data across the entire organization. These reference points can be as complex as customer interests and behaviors, or as simple as network configurations and authenticated users. In most cases, this dependency drives a requirement for data to be accessible at all times. Without that access, businesses and organizations can suffer massive disruption to their operations and consequently large financial impacts.

This sensitivity to loss of data access has spawned a multi-billion dollar industry, with losses estimated at over \$600 billion a year. Cyber criminals are becoming much more sophisticated and highly trained in data protection and security techniques in order to circumvent them. Incursions can be difficult to detect and, in many cases, the malicious users have lurked in an environment for weeks or even months prior to initiating an attack. The insight gained during that time allows attacks to target a large percentage of the environment and are calculated to achieve maximum disruption to the target. In many cases, they are aware of the need to disable or encrypt backup data as a first step in their attack. With no recourse for recovery, targets are more likely to pay for returned access to their data.

Unfortunately, even when ransoms are paid, there is no guarantee that access to all data can, or will be, restored. When ransoms are not paid, it typically takes weeks to identify all the infected systems and then find copies of backup data that have not been encrypted, deleted, or affected in some way.

Recent history has shown that an organization's likelihood of being affected by a cyber-attack is significantly higher than a disaster recovery (DR) event. In many cases current budget allocations do not reflect that reality. This threat has prompted many organizations to introduce solutions intended to maintain a copy of critical data in a secure environment, to ensure a safe repository to recover from. Dell EMC has introduced the PowerProtect Cyber Recovery Vault to address this exact use case.

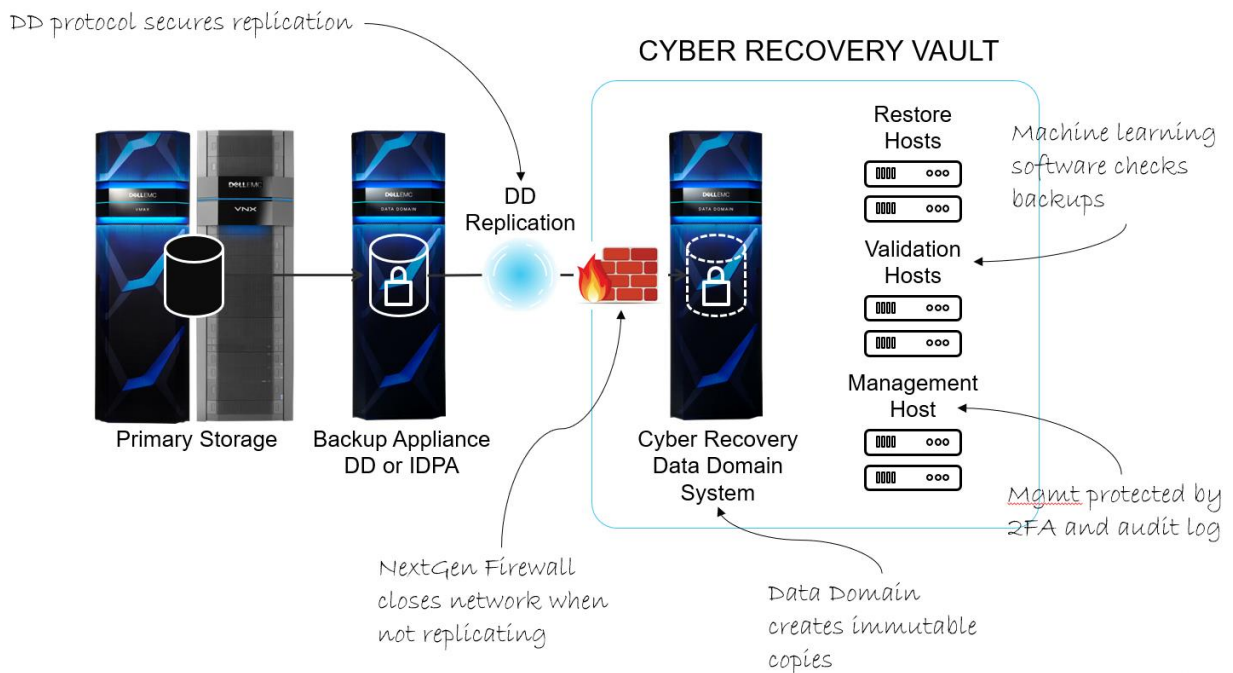
## 2. Dell EMC PowerProtect Cyber Recovery Vault

PowerProtect Cyber Recovery Vault automates the secure replication of data from a production Data Domain (DD) system to a Data Domain within a hardened vault location. The replication occurs over an air gapped replication link that is only open during the replication window. Data written to the vault leverages retention locked copies of the backups to provide data immutability. Further, once data is within the vault it is scanned and machine learning algorithms are applied to that data to automatically and quickly detect unusual behavior that is indicative of a cyberattack.

Cyber Recovery Vault Provides:

- Isolated, operational air gap for data vaulting
- Immutable copies of backup data
- Intelligent, secure, vault-based analytics with full content analysis
- Multi-layered security design protects against a full array of threats – including insiders
- Automated, orchestrated operation with a modern UI
- First Solution Provider member in the Sheltered Harbor Alliance program

### High Level Architecture



The base-level Cyber Recovery solution architecture consists of a pair of PowerProtect DD systems and the Cyber Recovery management host. In this base-level configuration, the Cyber Recovery software, which runs on the management host, enables and disables the replication Ethernet interface on the PowerProtect DD system in the Cyber Recovery vault to control the flow of data from the production environment to the vault environment.

## Key Components

- **Source (production) PowerProtect Data Domain system** — The source DD system contains the production data to be protected by the Cyber Recovery solution.
- **Destination (vault) PowerProtect Data Domain system** — The DD system in the vault is the replication target for the source PowerProtect DD system. This system can be a physical Data Domain system, or a Data Domain Virtual Edition when deployed in cloud environments.
- **Cyber Recovery software** — The Cyber Recovery software orchestrates synchronization, manages, and locks the multiple data copies that are stored on the Data Domain system in the Cyber Recovery vault, and orchestrates recovery. The software also governs the optional process of performing analytics on the data that is stored on the PowerProtect DD system in the Cyber Recovery vault using CyberSense.
- **MTree replication** — MTree replication is a Data Domain feature that copies unique data from the source Data Domain MTree to the Data Domain MTree in the Cyber Recovery vault.
- **Retention Lock (governance or compliance) software** — Data Domain Retention Lock software provides data immutability for a specified time. Retention Lock functionality is enabled on a per-MTree basis, and the retention time is set on a per-file basis. While not mandatory for Cyber Recovery, Retention Lock is strongly recommended as an additional cyber-resiliency measure.
- **Cyber Recovery management host** — The management host is where the Cyber Recovery software is installed. This server is installed in the vault environment.
- **Optional Recovery hosts** — The recovery host is a vault-environment component to which the backup application and data may be recovered. Vault environments may include multiple recovery hosts.
- **Analytics/indexing host** — Technically, the analytics/indexing host is optional, but is a highly recommended component in an effective vault environment. An analytics/indexing host with the data-analysis software that is installed provides direct integration between the Cyber Recovery software and the CyberSense software. Additional analytics/indexing hosts with different tools can also be used as needed.

## Optional Components to enhance security

- **Data Diodes** — A data diode from OWL Cyber Defense Solutions offers secure one-way communication from within the vault environment to the production environment for UDP Protocols such as SMTP and SNMP alerts.
- **ZeroTrust Networks** — ZeroTrust Networks within the Vault Environment are enabled using Unisys Stealth.
- **Firewalls** — Firewalls can be installed on the replication data path to ensure that only expected data traffic can traverse the secure link into the vault. The link must connect

directly to the Cyber Recovery vault PowerProtect DD system and not go through the Cyber Recovery vault switch.

### Key features of the Cyber Recovery software

The Cyber Recovery software, which runs in the vault environment, controls the replication interface on the Data Domain system in the Cyber Recovery vault as well as the Data Domain data copies in the vault. The software is built on a secure microservices architecture and provides the following key features:

- HTML5-based UI built on the Dell Clarity standard
- REST API implementation
- Command-line interface
- Informative dashboards showing system alerts, Cyber Recovery vault state, and other critical details
- Ability to transmit alerts through the SMTP to an environment outside the Cyber Recovery vault
- Cyber Recovery policy creation and management, including scheduling
- Recovery assistance and the ability to easily export data to a recovery host
- Automated Recovery options for the NetWorker and PowerProtect Data Manager applications
- Integration with the CyberSense software for detection of backup data that has been tampered with, including scheduling

## 3. Why is Dell EMC PowerProtect Cyber Recovery Vault a more effective solution?

### Cyber Recovery Air Gap

The air-gapped Cyber Recovery vault environment has both a physical and logical separation from the production environment. The separation reduces the attack surface of the Cyber Recovery vault. The base-level design for the vault network starts with the vault having its own network switching infrastructure. No intra vault communication is routable to any other environment. The only connectivity between the vault and another environment is as follows:

- Replication data link between the vault-environment and production-environment PowerProtect DD systems
- Optional dedicated link from the Cyber Recovery management host in the Cyber Recovery vault to the production network operations center or security operations center for events reporting

The Cyber Recovery software manages the replication link, and the connection is enabled only when new data must be ingested by the PowerProtect DD system in the Cyber Recovery vault. The Cyber Recovery software manages the link by enabling and disabling the replication port

and replication context on the PowerProtect DD system in the Cyber Recovery vault. Therefore, the replication link on the PowerProtect DD system in the Cyber Recovery vault uses its own unique Ethernet interface. The efficient replication enabled by Data Domain's unique deduplication capabilities ensure this link is only open and visible for the shortest possible window. As soon as the data transfer is complete, the network connection is not only closed, but the interface itself is disabled, thus reestablishing complete vault isolation via an operational air gap.

### **Why This Matters**

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. A technical validation performed by ESG has confirmed 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.

### **Full content analytics of data within the Vault**

Once data is replicated to the vault, CyberSense scans the backup image and generates analytics on that data. Analytics look inside the files and databases to uncover unusual behavior that is indicative of a cyberattack. This evidence includes file corruption, encryption of files or pages in a database, or deletions and creations.

The statistics are then analyzed using machine learning algorithms that have been trained on the latest ransomware threats to make a deterministic decision on whether the data has been attacked. If an attack has occurred and data corrupted, CyberSense delivers forensic tools to find the corrupt files, report on the user account that caused the corruption, so this account can be locked, and also will report on the application that made the changes to the file. With these forensic tools you can recover and diagnose a ransomware attack and replace corrupted files with the last good copy.

Initial scan of a backup image can detect an attack with up to 95% accuracy. Subsequent passes increase detection capabilities up to 99%. Together these solutions provide a secure and powerful solution against ransomware attacks. If an attack does get past the real time defenses, and corrupts files or databases, CyberSense can detect it quickly and within a backup cycle the last good copy of the data can be retrieved.

This rapid response enables business operations to continue without any interruption and cyberattacks to be thwarted quickly and painlessly.

### **Why This Matters**

An effective cyber protection strategy must be able to detect malicious behavior prior to valid backup retention periods expiring. CyberSense delivers this by looking inside the backup data to determine if there has been an attack. If needed, forensic tools are then used to find corrupt files, diagnose the attack vector, and identify a known good backup set. This ability to detect intrusion and then restore from the last good file minimizes business interruption.

## Ability to Recover following an attack

Recovering data from the vault in the event of an attack, or simply for testing, is critical. The Cyber Recovery Vault 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 these specifics, during a first pass analysis, on portions of the data set that may have been impacted. Further, 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 be the victim 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 the ability to recover business-critical systems as quickly and efficiently as possible after a cyber incident.

## Deployment options for a Cyber Recovery Vault

There are two common ways being leverage to deploy Cyber Vaults today.

1. Many organizations are deploying Cyber Vaults in one of their own physical data centers, or securing co-lo space in a third part data center, to house this infrastructure. In these scenarios, the goal would be to provide the most secure infrastructure possible with the only communication into the vault occurring via a one-way air gapped replication link. Although it is becoming more common to also allow secure log shipping out of the environment for management and faster alerting.
2. Other organizations do not want the vault within their own data centers or may not want to provision a secondary data center to contain the vault. Increasingly those organizations are looking to cloud infrastructure providers such as Microsoft Azure to house their Cyber Recovery Vault.



## 4. Dell PowerProtect Cyber Recovery Vault deployed on Microsoft Azure

As a cloud service provider Azure has a proven track record providing a secure environment for client workloads. Enhanced security combined with a flexible environment to enable fast deployment along with dynamic expansion or contraction of vault resources deliver flexibility and cost management.

The Cyber Recovery software manages a virtual air gap between a production environment and the Cyber Recovery vault. It disables replication links and replication ports on the production Data Domain system when Cyber Recovery policies are idle. The software enables and disables access to both a private subnet and Data Domain Virtual Edition (DDVE) in the Cyber Recovery vault, which are installed during the solution deployment, through security groups and ACLs.

When a policy runs, the Cyber Recovery software enables the flow of data into the Cyber Recovery vault by enabling both the replication link and the replication port of the Data Domain system. When a policy finishes synchronizing data into the Cyber Recovery vault using the replication link, the Cyber Recovery software disables the replication link. Also, when all policies no longer use a specific Data Domain port to synchronize data into the Cyber Recovery vault, the Cyber Recovery software disables the port by bringing down the interface.

Azure provides Virtual Network (VNet) security mechanisms that provide additional security measures for the Cyber Recovery vault:

- Security groups, which protect the instances deployed in the VNet
- Network ACLs

The Cyber Recovery software enables and disables access to a private subnet through a network access control list (network ACL) and enables and disables access to an instance through security groups.

### PowerProtect Cyber Recovery for Azure

The Cyber Recovery software is made available as an Azure Virtual Machine image. To deploy the Cyber Recovery software to an Azure Virtual Machine instance in a Virtual Network (VNet), use an Azure Resource Manager (ARM) template.

The CloudFormation template creates:

- The Cyber Recovery VNet—The VNet includes all the components required for the Cyber Recovery solution.
- Two subnets—The two private subnets include:
  - An Azure jump host on one subnet
  - The Cyber Recovery management host and DDVE on the other subnet

The production workstation cannot access the Cyber Recovery management host directly. The Windows-based jump host is available in the VNet to access the Cyber Recovery and DDVE instances. The management path is through the jump host.

- Network access control lists (ACLs)—The ACLs provide a layer of security for the VNet that act as a virtual firewall for controlling traffic in and out of the subnets.
- A security group for each instance—The security group protects the instance by acting as a virtual firewall to control inbound and outbound traffic.
- VNet endpoints—The VNet endpoints enable private connections between the VNet and supported Azure services.
- Identity and Access Management (IAM) roles—Along with the VNet endpoints, the roles provide access to Azure services for specific instances.

The CloudFormation template also deploys an Azure jump host. The Windows-based jump host is available in the VNet to access the Cyber Recovery and DDVE instances. The management path is through the jump host.

Back up data is stored in Azure Blob Storage buckets with a high level of deduplication.

The Cyber Recovery deployment using an ARM template does not include a VPN. We strongly recommend that clients:

- Set up a VPN.
- Use a VPN gateway or Azure ExpressRoute to access the jump host.

## PowerProtect Cyber Recovery and DDVE

To function on Azure, the Cyber Recovery software requires that Data Domain Virtual Edition (DDVE) is also installed on the Azure VNet. The Cyber Recovery solution deployment on Azure installs DDVE.

DDVE is a software-only protection storage appliance: a virtual deduplication appliance that provides data protection for entry, enterprise, and service-provider environments.

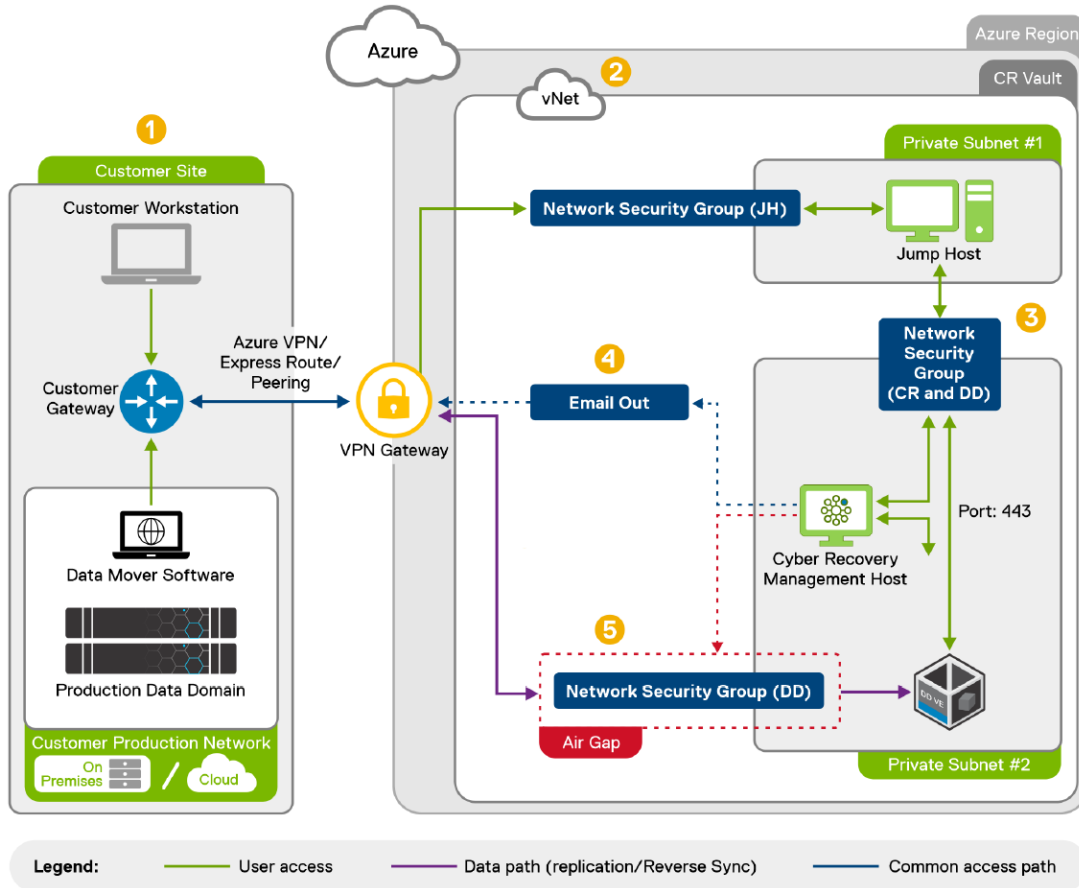
Leveraging Data Domain Virtual Edition within Azure delivers a cost-effective solution reducing TCO in the following ways:

- Industry leading deduplication to decrease cloud storage costs
- All data replicated into the vault is deduplicated prior to replication to reduce bandwidth requirements and time to replicate
- DDVE can be provisioned in as little as 1TB increments so you only need to provision the capacity required to support the vault

## Architecture overview

The basic Cyber Recovery solution on Microsoft Azure architecture includes a single region, a single Virtual Network (VNet), and a single availability zone (AZ).

The following figure represents the architecture. The right side of the figure shows the Azure resources that define the Cyber Recovery vault architecture:



1. The production environment can be on premises or also deployed on Azure. The workstation at the production site enables you to connect to the jump host, which is in a private subnet in the VNet. The jump host provides access to the Cyber Recovery management host and the DDVE management console. For additional security, the workstation has a limited IP range.

2. The ARM template deploys all the components that the Cyber Recovery solution requires in the VNet on Azure. The template creates two private subnets: A private subnet that includes the jump host and a private subnet that includes the Cyber Recovery management host and DDVE. It also configures security groups, Access Control Lists (ACLs), inbound and outbound rules, and so on.

3. The network Security Groups allow access between:

- The production workstation and the jump host subnet
  - The private subnets that include the jump host, the Cyber Recovery management host, DDVE, and the other components that make up the Cyber Recovery vault
4. The Cyber Recovery email capability provides one-way email from the Cyber Recovery management host. You can also use third-party email services at an additional cost.
6. The Cyber Recovery software automatically enables and disables the air gap, which uses Azure security features for additional security.
7. The Data Domain system and DDVE use bi-directional data communication.

Sensitive Cyber Recovery data, such as passwords, is encrypted and stored in a lockbox. For more information about Cyber Recovery security, see the PowerProtect Cyber Recovery Security Configuration Guide at Dell Online Support. When deployed to Azure, the Cyber Recovery lockbox is located in a secure Managed Disk.

Backup data is stored in Blob Storage, and the backup metadata is stored on a Managed Disk.

### **Assigned resources in the Cyber Recovery Azure virtual network (VNet)**

The Cyber Recovery solution deployment assigns system resources for the instances it creates in the Cyber Recovery Azure virtual network (VNet).

The system resources for the three Cyber Recovery solution instances include:

- For the Cyber Recovery management host:
  - D4ds v4 instance type
  - SUSE Linux Enterprise Server 12
- For the DDVE compute:
  - D4ds v4 instance type (minimum)
  - DDVE DDOS 7.8.0.0
- For the DDVE storage:
  - Azure Blob Storage
  - Standard SSD LRS
- For the jump host:
  - Standard D2 v3 instance type
  - Microsoft Windows Server 2019 Benchmark - Level 2

## Cyber Recovery on Azure limitations and unsupported features

Before you deploy the Cyber Recovery solution on Azure, review the following limitations and unsupported features.

Note the following:

- The maximum number of supported policies is 32.
- Azure does not provide an integrated email service. Use the Cyber Recovery email capability. You can use third-party email services at an additional cost.
- The CyberSense feature and Sheltered Harbor deployments are not supported currently.

## 5. ESTI Consulting Services

ESTI are experts in IT Consulting with a focus on data management and protection solutions and extensive expertise around the Dell PowerProtect data protection portfolio. We also have years of experience helping clients understand their application and data assets to determine which are prime targets for movement to the cloud. And finally, we work hand in hand with those clients to support migration efforts to Microsoft Azure and any modernization of applications required to support those efforts.

With Respect to a Cyber Recovery Vault (CRV), ESTI will work with you to help articulate the differences between DR and CRV. The goal of CRV and data targeted for protection are different than the requirements of DR. ESTI Business Analysts will work with your team to understand business processes and document your application catalog and infrastructure. We then help you define your cyber recovery priorities, and identify valuable business processes, key applications, and ultimately key data to be protected.

## 6. Summary

Dell EMC PowerProtect Cyber Recovery Vault, when provisioned on Microsoft Azure, provides a highly effective vault solution. The Vault can be implemented quickly to protect your data from vulnerability, and your organization from lengthy business disruption resulting from a cyber or insider attack. All this occurs within a secure Azure Virtual Network that offers an agile and flexible deployment to manage costs. ESTI Consulting Services has the data protection expertise to identify data requiring placement into a vault, stand up a Cyber Recovery Vault within Azure, and integrate the vault into your production environment.

## 7. References and Additional Documentation

James Lewis (2018): Economic Impact of Cybercrime — No Slowing Down  
<https://csis-website-prod.s3.amazonaws.com/s3fs-public/publication/economic-impact-cybercrime.pdf>

Dell Technologies Solutions (2022): Dell EMC PowerProtect Cyber Recovery Solutions Guide  
<https://www.delltechnologies.com/asset/en-us/products/data-protection/technical-support/h17670-cyber-recovery-sg.pdf>

Vinny Choinski, Christophe Bertrand (2020): Protecting Critical Data from Cyber Threats Such as Ransomware with a Comprehensive Digital Vault Solution  
<https://www.delltechnologies.com/resources/en-us/asset/analyst-reports/products/data-protection/esg-cyber-recovery-tech-validation-report.pdf>

Dell Technologies Solutions (2022): Dell EMC PowerProtect Cyber Recovery 19.12 Azure Deployment Guide