DELL Technologies

Whitepaper

Recover Your Critical Data in Case of Cyberattack

Leverage Dell PowerProtect Cyber Recovery

Recovering from a Cyber Event

The recovery process from a serious encryption or data destruction attack is unique to each event. Incident Response Teams (IRT) will evaluate numerous factors, including the scope of the attack, business requirements to return to service, uncertainties about containment and eradication of the threat actor, legal and regulatory concerns, etc. to select the best path and method of recovery. Because of this process, enabling a variety of recovery options is useful.

This paper discusses the steps needed to properly and safely respond to and recover from different ransomware and data destruction attacks, ranging from small in scope to large and severe. In addition, we discuss how Dell's cyber resilience capabilities, including an isolated cyber vault, enable a variety of different recovery options, along with a discussion of the benefits for each.

July 2025

Revisions

Date	Description
June 2022	Initial release
July 2025	Additional recovery options and updated images

Table of Contents

Responding to a Cyber Incident

Organizations can choose from a variety of incident response models and frameworks and should select the one that best suits their needs.

The basic workflow depicted below is adapted from the NIST 800-61 SP and has proven useful in helping organizations to understand the steps necessary in recovery from a cyber-attack: [1]

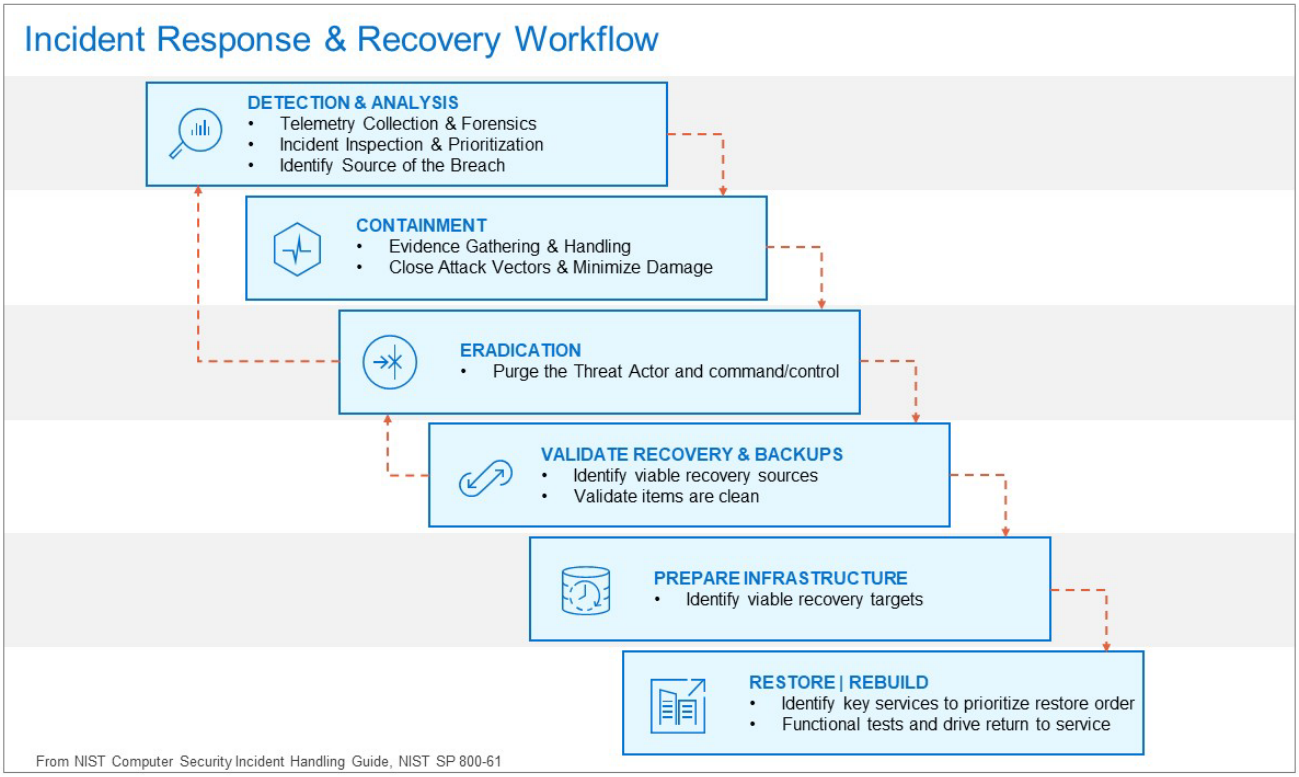

Detection & Analysis. The primary goal in the initial step is to determine whether an adverse incident has occurred and to begin to understand its impact and scope. As described in NIST SP 800-61 r2:

> When the team believes that an incident has occurred, the team should rapidly perform an initial analysis to determine the incident's scope, such as which networks, systems, or applications are affected; who or what originated the incident; and how the incident is occurring (e.g., what tools or attack methods are being used, what vulnerabilities are being exploited). The initial analysis should provide enough information for the team to prioritize subsequent activities, such as containment of the incident and deeper analysis of the effects of the incident.

Another critical component of this phase is reporting the incident to appropriate personnel and external organizations and beginning to answer key questions:

- What was impacted?
- How much damage was done?
- Where is the source of the attack?
- When did the attack begin?

Containment. This stage focuses on stopping or at least limiting any further threat actor activity in the environment. It may require shutting down or disconnecting systems, cutting access from outside the data center (or public cloud environment) or even unplugging or disabling switches and other connections.

Often, this stage also includes the creation of forensically sound images for key systems, which will aid in both understanding the root cause and scope of the attack and serve as evidence for use in later legal or regulatory proceedings.

Eradication. After containment, the threat actor's ability to persist in or return to the environment must be eliminated. This can be a complex task. The MITRE ATT&CK Framework (attack.mitre.org) currently identifies twenty techniques used by threat actors to maintain persistence, including:

- Create Account
- Modifying Authentication Process
- Valid Accounts

In addition to persistence capabilities, if the initial access vector itself has not been remediated – such as an unpatched vulnerability – the threat actor can simply return to the environment at any time.

Eradication is critical. Failure to eliminate threat actor access will often result in the attack re-starting after recovery has begun. This not only sets the organization back to the recovery starting point, losing valuable time, but also erodes credibility with customers and partners. Attackers also do a more thorough job of destruction when given a second chance.

Validate Recovery and Backups. As the actual restore / recovery process begins, clean sources of data must be identified. A backup is often a primary source for recovery if it survives the attack and the backup infrastructure is available or can be quickly rebuilt to enable access. An isolated vault (which is described more completely below) is an even more valuable source. In some cases, other sources of data may be used in recovery: snaps from primary storage, offline tape backups and even data decrypted by paying a ransom or exploiting mistakes by threat actors.

A frequent challenge in using any of these sources relates to access and availability. Some potential scenarios include:

- Active Directory is a frequent target, and may need to first be rebuilt or revalidated before the IRT can authenticate and access data on systems using it for authentication
- Primary storage platforms may have been attacked or over capacity due to threat actors or snaps being automatically taken to keep up with rapid changes to data (e.g., encryption)
- Forensic image work may need to be completed before changes to the systems are permitted
- The systems may reside on a network or network segment that is untrusted and subject to further attack

Once accessed, data may need to be scanned or scrubbed to validate useability. At this stage, the best recovery techniques can be evaluated, along with prioritizing and sequencing the recovery of specific systems and applications. Key questions to be answered in this stage include:

- What source(s) contains a last known good copy of data that is accessible?
- How can the data be validated as useable?
- What is the order in which applications should be restored?

Prepare Infrastructure. The clean data (including applications) next requires a location to be restored and then run. There are several viable options to prepare the infrastructure. Depending on the attack scope and sophistication, risk tolerance and other factors, systems may be wiped to the OS level, or even to the firmware level. New, borrowed, or re-purposed equipment is often brought in for use in the recovery. Public or private clouds may be an option in some cases. If necessary, additional monitoring may be implemented to rule out threat actor activity on restored infrastructure.

A question often asked at this stage is why all existing infrastructure cannot just be wiped and re-used. That may be possible. But in many cases, existing infrastructure may hold encrypted data which may be held to enable a recovery – as a last resort – by paying the ransom. Alternatively, the storage may have partial copies of data that are being slowly restored or repaired, making the platform unavailable for new use.

Restore and/or Rebuild. Finally, clean data is restored to clean hosts and infrastructure. As noted, this step may take various forms including a bare metal rebuild, restoring from backups, rebooting virtual machines, etc. The best choices depend upon an analysis of the attack, an understanding of what has been impacted, and the organization's risk tolerances.

Resilient Recovery with Dell PowerProtect Cyber Recovery

Backup data stored on a Data Domain platform, with immutability and aligned to zero trust principles, provides a strong base for any recovery operation. The additional option of incorporating a secure, isolated cyber vaulting capability – such as PowerProtect Cyber Recovery (PPCR) – will further streamline and enhance the recovery process, improving efficiency and recovery times.
To be useful in a recovery scenario, a cyber vault must have three characteristics:

1. **Isolation.** The components of the data vault must be physically and logically isolated. Logical isolation has similarities to an air-gapped network, except for limited connectivity for replication of data into the vault. Key to this capability -and unique to PPCR - is that control type traffic (SSH, HTTP/S, etc.) is never allowed in the vault environment from any outside source, including production. This isolation ensures that there is no risk of threat actor activity, control, or persistence in the environment.

2. **Immutability.** All data written to the data vault must be secured in a manner that electronically prohibits deletion or modification until the expiration of the retention period, which is typically a couple of weeks. Immutability provides additional protection against an insider threat and ensures that data which has been validated has not been changed.

3. **Intelligence.** Data in the vault should be analyzed or interrogated in a manner that ensures the data has not been manipulated or corrupted. This capability validates that data in the vault is in a known "good" state unless flagged otherwise. Importantly, this means the data can be used, immediately, to begin recovery operations.

Recovery Options

Coupled with the Data Domain base resilience platform, PPCR enables many different recovery options from which to select. This enables the IRT, after evaluating the scope and extent of the attack, to recover the organization in the safest, most efficient manner.

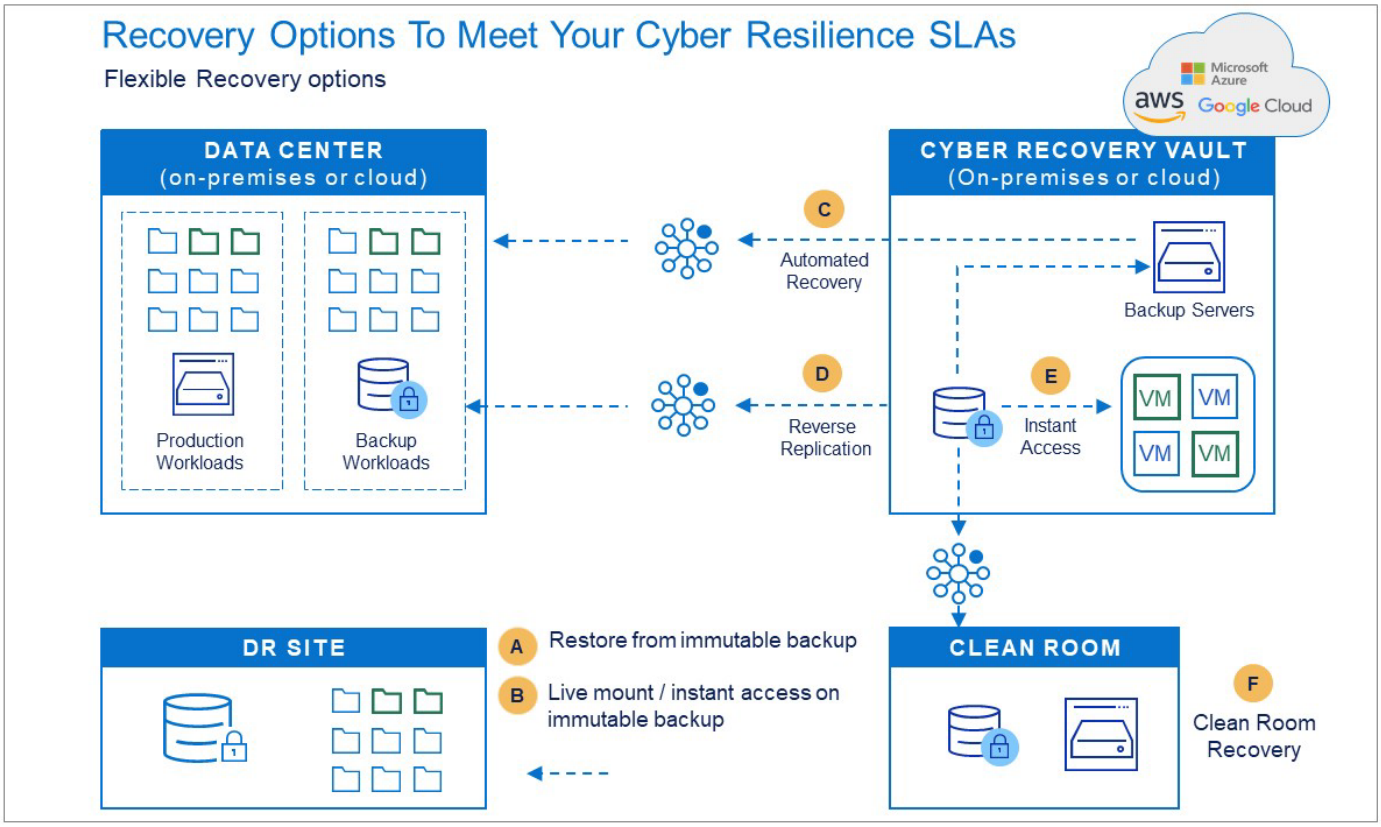

Option A – Restore from An Immutable Backup

In situations where the attack is minor or was quickly contained, the best option may be to restore a small number of files, VMs or servers directly from the backup environment. The immutable storage capability of the Data Domain helps to prevent threat actors from modifying, encrypting or deleting backups or their catalogs.

Benefits: Restoration can be fast and simple

Limitations: Probably limited to less sophisticated attacks, or attacks that are stopped in the initial stages

Option B – Live Mount / Instant Access

Similar to Option A, the best recovery path from a minor attack may be to quickly boot impacted virtual machines directly from the Data Domain. These VMs can be started quickly, and then vMotioned back into production as the hypervisor stack is fully recovered and validated.

Benefits: Extremely fast restore, effectively leverages virtualization

Limitations: Dependent on VMWare capabilities; only a limited number of VMs can actually be run simultaneously (although many can be mounted)

Option C – Automated Recovery from the Vault

In scenarios where the production environment may not be immediately accessible or quickly validated for use, or the backup environment has been impacted, the cyber vault provides a known good copy of data. Normal vault configurations include an instance of the backup software (such as PowerProtect Data Manager), which can

be used to access and recover needed files (data, applications, configurations, etc.) from the vault and restored to the appropriate (or new) production servers. This process can be automated.

Benefits: Recoveries are very secure because they utilize vaulted data. In a smaller recovery, the recovery can be very fast

Limitations: Recoveries involving large amounts of data will take longer

Option D – Reverse Replication from the Vault

Some scenarios may benefit from having a known, clean copy of data on what was previously the production-based backup target. For example, if a significant amount of automation work has been created for recovery, source data may be required to be staged on that device. If the target has been impacted by the attack, it can be wiped / reset, and the vault source used to restore its full data set. Note that this extra step will take longer than a direct restore (Option C) because the data must first be replicated; but that replication process can be done independently from and in parallel to other recovery operations.

Benefits: Very efficient in certain recovery scenarios

Limitations: Reverse replication path should be planned to ensure the highest potential speeds. This will not be the bottleneck for recovery, and its impact can be limited when done in parallel with other recovery operations

Option E – Instant Access from the Vault

Like Option B, VMs in the vault can be mounted and run via the instant access capability. This option is particularly useful if the production backup environment has been compromised or cannot be immediately validated. Note that this option may require some pre-planning to enable network connections and possible changes to IP addresses, all of which can be incorporated into run books or as part of a larger automation project.

Benefits: Very efficient and fast recovery

Limitations: Pre-planning is necessary to identify specific VMs to be recovered; and infrastructure requirements to run from the vault must also be identified

Option F – Clean Rooms

Clean Rooms represent otherwise unused, isolated infrastructure and are primarily used for testing, data sanitization, validation, and application recovery to expedite the recovery process.

- *Option F1*- In this scenario, clean room infrastructure is sized by the organization based upon the largest application to be recovered. Once the integrity of the application has been established in the clean room, e.g., the applications are shifted back into the production environment and the IRT moves to recover the next application in the clean room. This represents the most common process in which the incident response teams can ensure all data is clean before it is recovered back into production.

  Benefits: In a severe attack, a clean room can save days to weeks during the recovery process

  Limitations: An investment in infrastructure specific to the purpose is necessary

- *Option F2* – This option is used by organizations desiring to recover a subset of most critical application(s) and make them accessible right away. In this scenario, recovery of the application is again completed in the clean room; but the application is run from the clean room environment after being promoted as a production environment. This will take additional pre-planning for networking, connectivity, etc. but can significantly shorten the time to recover.

  Benefits: The only method that can ensure a recovery without dependencies on other, unrelated processes

  Limitations: Requires substantial maturity and significant investment

Minimum Viable Organization Option

The vault also enables a very sophisticated form of recovery, often referenced as a minimum viable business or organization. A significant amount of pre-planning and maturity is required, but done well, this process can enable restoration of critical operations within hours to days.
This requires:

- A full understanding of the services most critical to the organization
- Identification of the applications and related dependencies required to deliver those services
- Pre-staged equipment, sized to enable the restoration and running of the above-referenced applications

As noted, this is a very mature and potentially expensive option. The figure below covers the different architecture options available to accelerate the recovery process. However, it is the only method to ensure that recovery operations can begin immediately after an attack – without waiting for eradication, containment, forensics, validation, etc.

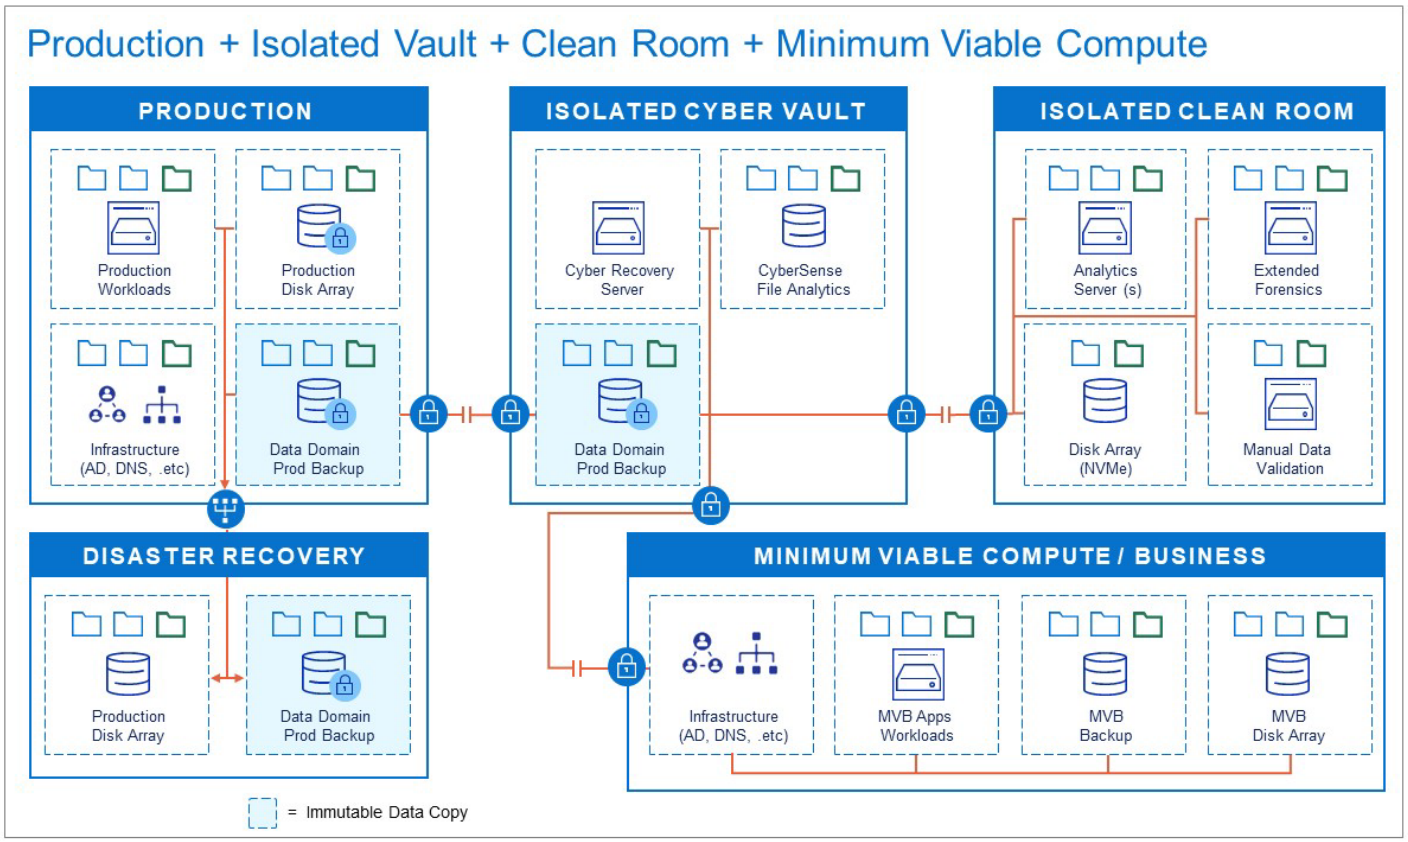

Conclusion

Dell PowerProtect provides the industry's most effective recovery solution against common attack vectors, including dormant malware, data wiping and locking, data corruption, insider attacks, and the destruction of backup and storage assets.

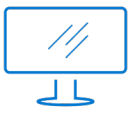

Learn more about Dell PowerProtect Cyber Recovery

Contact a Dell Technologies Expert

Demo Dell PowerProtect Cyber Recovery

Join the conversation with #PowerProtect

[1] NIST SP 800-61r3 Incident Response Recommendations and Considerations for Cyber Risk Management (April 2025).

DELL Technologies