druva

The Ransomware Survival Guide

Evaluate risks, costs, and innovative best practices to safeguard your data from cyber crime.

Introduction

Ransomware protection can be overwhelming.
Every media outlet leads with the devastation and omnipresence of ransomware. Every vendor touts their ransomware "solution" as if such a complex problem could be solved with just one product. Every CEO and board member demands a comprehensive ransomware strategy… within budget, of course. Where do you start?

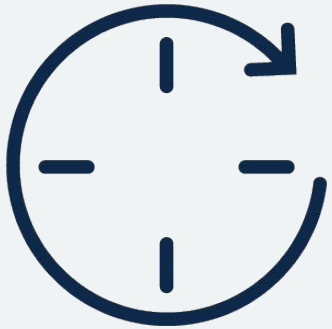

This guide will help you understand how to survive a ransomware attack.
We will share the requirements, best practices, and how to measure and share your progress. We will begin with the foundations, then explore advanced recovery services, and finally explain how to operationalize everything. After reading this guide, you will be able to better define the data protection component of your ransomware recovery strategy and effectively implement tactics that will get you back to normal faster after an attack.

A Strong Foundation

There's no need to be intimidated by the noise around ransomware recovery. The foundations of data protection for ransomware recovery are the same as the foundations for traditional recovery.

First, get a reliable third-party cloud data protection service in place for all of your data, and ensure that the protection storage and protection software cannot be compromised.

Second, keep the architecture focused on recovery performance and flexibility, because ransomware is evolving so quickly that premature feature optimizations will limit your ability to respond.

Finally, test your recoveries in as many ways as you can.

Data protection for ransomware recovery is still data protection. Ransomware provides an excellent opportunity to refocus your organization on data protection best practices, before moving to more complex functionality.

druva

Protecting Your Data

After an attack, your organization might take many steps, but they will begin and end with recovery. In the beginning, you will need to answer: "Can we recover the affected data in time, or will we have to pay the ransom?" In the end, you will need to deliver clean, unencrypted data.

Since cyber criminals know the importance of your backups, they will target them for corruption or deletion. Therefore, not only do you need to protect the data, you need to protect your copies. There are four requirements for data protection in the age of ransomware:

- **Reliable, resilient backups** — You should target a 99+ percent successful backup rate. More importantly, the backups should be resilient and durable, so you can target a 99+ percent successful recovery rate. If your backups are not working, ransomware does not even need to corrupt or delete them.
- **Unmodifiable backups** — Ransomware will try to corrupt or delete your backups. There should be no way for a process/server in your environment to directly access the backup storage — e.g. via NFS, SMB, S3, or local file protocols — because ransomware will exploit the link. Even "root only" accesses are a vulnerability because cyber criminals can gain root access to your systems.
- **Non-deletable backups** — Ransomware will also try to gain control of your backup software and delete all the backups. It can manually trigger deletes, reduce the retention period to force automatic deletions, and alter backup schedules so no new backups are created. Therefore, you need a system that will prevent backup deletion, even by an administrator, because administrator accounts can be compromised.
- **Back up everything** — Ransomware usually enters via end-user devices, but attacks any data source — SaaS applications, cloud applications, filer servers, VMs, or even databases. Therefore, you need to protect everything with the same resilient, unmodifiable, non-deletable approach, otherwise your recovery may be incomplete or missing critical data.

Since ransomware attacks SaaS applications, cloud workloads, data centers, and endpoints, you need a centralized backup solution that can protect all your data.

Architectural Flexibility

Every product has a unique ransomware protection feature, but it is impossible to create an architecture flexible enough to incorporate every differentiated approach. Furthermore, ransomware itself will continue to evolve, so you want a simple enough architecture that avoids becoming brittle. Therefore, you will need to decide where you want to build in architectural flexibility for the future.

There are three areas in which architectural flexibility will matter:

- **Scalable recovery** — Ransomware attacks try to compromise as much data as possible in as short a time as possible. Therefore, organizations will need to recover data both successfully and quickly. Since **protection** infrastructure is rarely sized for large-scale recoveries, it will become increasingly important to design for the flexibility to recover at scale.
- **Multi-cloud recovery** — It can take days or weeks to clean out ransomware from an environment. Organizations that can recover data to alternate locations can restart critical applications even while quarantining infected areas. Instead of losing revenue and customer confidence, the business can continue to run.
- **Ability to evolve** — Ransomware is constantly evolving, and so is your data environment. Therefore, you want an architecture that enables frequent upgrades of the protection environment. Threat actors can easily gain access to the backup environment via vulnerabilities that were not automatically patched; in fact, 42 percent of vulnerabilities are exploited after a patch has already been released. You cannot afford to use 18-month-old protection software any more than your security team could afford to use 18-month-old anti-virus signatures.

When it comes to the foundations of data protection, follow the aphorism: "It's not about backup… it's about recovery." You will need to recover more data faster. You will need to recover to new environments. You will need to recover workloads that don't even exist today.

The fight against ransomware is just beginning. Architectural flexibility is more important than any feature.

Testing Data Recoveries

The best way to protect your organization from a ransomware attack is to test your data recoveries every week.

Teams worry about testing realistic "ransomware recovery," but if you can't run a recovery, you can't run a ransomware recovery. While ransomware recoveries create additional requirements, which we will cover in the last part of the guide, the fundamental requirement is — you should be able to run a successful recovery 99+ percent of the time. You will only hit that mark if you test in the following ways:

- **Test recovery in different environments** — Depending on the scope of the attack and the urgency of the restore, you may need to recover your data and applications to an alternate environment. No matter how portable your applications are, you want the first cross-environment restores to be tests so you can work out performance, security and network settings, application dependencies, and the unknown unknowns unique to your organization.
- **Test recovering different workloads** — Most restore tests tend to look something like: "Recover a few VMs, some files, and a tablespace." A ransomware recovery could require restoring SharePoint Online, a NAS share, 40 laptops, and dozens of VMs. You need to know both the functional and performance limitations of restoring different workloads, or you won't be able to know for certain whether you can recover the affected data in time or have to pay the ransom.
- **Test them with different people** — You want your entire staff to be able to run recoveries because recovering at scale does not just require technology — it requires people. The only way to be comfortable under pressure is to practice.

Practice makes perfect. Ransomware recovery is the motivation you need to practice.

Unique Challenges

Ransomware recovery extends the foundations of traditional recovery to address the unique challenges brought about by a cyber attack.

First, unlike a traditional failure, the ransomware protection solution needs to help identify, analyze, and assess the damage of a ransomware attack. **Second,** ransomware recovery needs to be a two-staged process to ensure that you are eliminating malware and restoring the latest good version of your data. Finally, since ransomware recoveries tend to be extensive, you need a validated plan and scalable infrastructure to meet the timelines of the business.

Rapid recovery from a ransomware attack can mean the difference between your business surviving the attack or collapsing under lost revenue and customer confidence. **You need to be ready.**

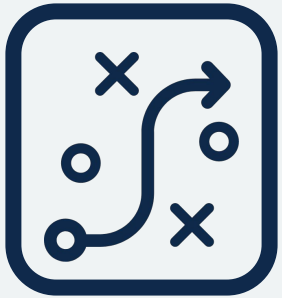

A note on detection

Your primary environment security solutions are your most powerful tool when it comes to detecting a ransomware attack. Your data protection solution should help your security team confirm the details and timeframe of an attack by identifying anomalies, and sending alerts to your security management. Detecting a ransomware attack is complicated, but information about unusual behavior and activity in your backup environment can make detection faster and more reliable.

Contain, Identify, and Recover

Ransomware can lie dormant within your environment for months, infiltrating your systems. When everything is "in place," the cyber criminals launch a massive coordinated attack. They will first target the backups, then encrypt as much production data as possible so that recovery will seem like an insurmountable task. The first way to optimize your recovery is to ensure your backup solution is coordinated with your primary environment security tools. This orchestration can help you quickly contain an attack, assess the damage as quickly as possible, and recover complete and clean data.

The first step after a ransomware attack is to stop the spread of malware. Backing up then restoring corrupted files can take you back to square one by reintroducing ransomware to your systems. Your backup solution should have the ability to automatically quarantine affected resources to avoid reinfection while you work to understand the scope of your ransomware attack.

Second, you need to figure out where the attack started and how it spread. Understanding the time frame and details of the attack is vital to identifying the correct data to recover in the final step of this process. Collaboration between IT and security teams is key in this step of recovery. Your backups should provide historical information to your forensic analysis tools to expedite the process. Historical logs can be useful for tracking the progress of the malware, and catalog searches can identify when/where malware files arrived onto OneDrive, a VM, or a NAS share.

Other stakeholders also need the information that comes from this analysis. The legal and executive teams need to understand what data is affected, so they can explain the magnitude of the damage to auditors, board members, and customers. Backups can provide data for risk assessment and remediation.

Once you have stopped the spread of ransomware and clearly understand the scope of the attack, you are ready to recover your data. There are two primary concerns when it comes to recovery, restoring only clean data and restoring it quickly. Let's talk about best practices for addressing both concerns.

Validate Before Recovering

There are two rules to recovering from a ransomware attack:

- **Do not** restore malware back into the environment
- **Do** restore the most recent good version of data

You must remove ransomware prior to restoring any system into the environment. Ideally, your protection vendor should scan for and remove malware before the restore. Regardless, we recommend that you also run your own malware scan. Many customers bring up the restored data in an isolated environment, run their own scans, and then proceed to bring the data into production.

With the cloud, running a preliminary restore does not have to be expensive and slow. Since the cloud spins up resources on-demand, you do not need separate "standby" resources for the first stage of ransomware data recovery. Additionally, recovery to a cloud instance (especially from a backup in the cloud) can be extremely fast. Finally, recovery is fast and on-demand, so you can set up the restored instances while still in the "analyze" and "assess" phases.

You should recover the most recent good version of your data with a combination of analytics and self-service. First, determine if your protection vendor can detect anomalies — this can immediately eliminate corrupted backups. Second, look at the distribution of file types across different backups and discard those with unusual backup types. Third, even after the recovery, users should be able to rapidly extract files from older backups with self-service restores.

Automation can also greatly reduce the manual effort necessary to accomplish this. AI technology has made it possible to identify the most recent clean version of every file or data set across the entire time frame of an attack, then compile them into a single snapshot so you can recover clean and complete data immediately.

The entire business is watching, so you only want to restore to production once. Use analytics, built-in malware scans, and test restores to ensure that you are ready.

Recover Quickly

Once you are prepared, you want the ransomware recovery to run as quickly as possible. Most recovery performance comes from preparation, so you need to prepare.

The three key steps to recovery performance are:

1. **Prioritize** — Under stress, every business struggles to identify which applications and infrastructure should be recovered first. Therefore, create a recovery plan ahead of time. The business can identify what matters, so when it comes time to recover, you just have to execute.
2. **Recover applications, not infrastructure** — The business cares about applications, and ransomware may affect some components of an application but not others. In addition, you can further break an application by restoring components to a previous point in time. It is critical to test restores to validate application dependencies so you can recover the application when the time comes.
3. **Cloud scalability** — Most on-premises environments are not built for large-scale recoveries, and they can bottleneck on protection appliances, network, and even the production target. The cloud can enable on-demand scale in all three dimensions: storage, compute, and network. Recover from the ransomware attack first and repatriate workloads to your data center at your own pace.

There is no magic formula to rapid recovery, but the cloud is a key ingredient to success. If you are prepared, your recovery will be focused, successful, and run at the scale of your business.

Operationalizing Ransomware Protection

Ransomware protection has to be operationalized because it cannot be solved by a single product, process, or person alone. Ransomware attacks are constantly evolving, so ransomware protection must evolve to meet them. The attackers try to gain control of the environment, so organizations must adopt zero-trust security across their entire backup environment — hardware, software, and the cloud.

Perhaps most importantly, ransomware protection must be cost-optimized and aligned with overall business objectives. By addressing ransomware as an organizational and operational challenge, instead of a technology challenge, you can stop reacting to threats and take the fight to the cyber criminals.

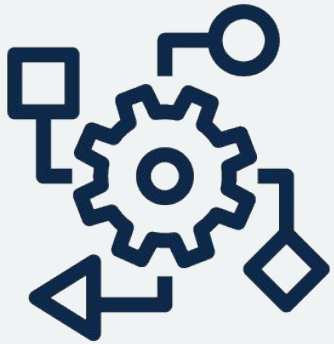

Optimize The Risk/Cost Equation

Ransomware is a devastating threat, but companies are not handing their IT departments blank checks to address the problem. In enterprise data protection and security, there is always a balance between cost and risk.

Optimize your infrastructure cost by asking three questions:

- **How many backup copies do I need** for each type of data? In some instances, you may just want one air-gapped backup copy. In others, you may want a local copy and a remote air-gapped copy.
- **Where will I store the air-gapped copies** — tape, additional backup appliance, or the cloud?
- **How much network bandwidth will I need** to make and update the air-gapped copies?

Most organizations do not always want two backup copies because they cannot afford the additional storage and networking costs. Therefore, they adopt a model that, by default, creates network-efficient ransomware-protected backups while allowing them to selectively create local backups.

Second, optimize the administrative cost. Once again, ask yourself three questions:

- **How many people will I need to manage the air-gapped network** — both the security and the bandwidth provisioning?
- **How many people will I need to manage the air-gapped storage** — tape, backup appliance, or the cloud?
- **How many people will I need to ensure that the environment is secure** — security patches, vendor best practices, immutable storage management, etc.?

These are not challenges you can simply "solve" and then ignore. The dynamic nature of the environment and the threats mean that this will require ongoing investment. Since most organizations do not want to add IT infrastructure administrators, you will need to slow down other projects to support ransomware protection. As a result, most organizations adopt a solution with built-in ransomware protection.

Find a solution that minimizes infrastructure and management costs, so you can reduce risk within your budget.

Implement Zero-Trust

Zero-trust has become a security foundation. Zero-trust means your organization doesn't allow access to systems for data for anybody or anything that is not authenticated and authorized. It's also important to monitor the data and users that have passed through. In recent attacks, we've seen ransomware take over production systems, emails, and even phone systems. When ransomware can compromise your entire environment, zero-trust is the only possible response.

Ransomware targets backups because cyber criminals know backups are your best defense against paying the ransom. If your backups are on-premises, ransomware tries to encrypt, corrupt, or delete them. If you have offsite backups, the cyber criminals try to gain control of your backup system to delete them.

To protect against ransomware deleting your backups, you need to meet, at a minimum, the following zero-trust requirements:

- Multi-factor authentication **(MFA)** — With MFA, the cyber criminals will need to compromise multiple components of your infrastructure, not just one.
- **Four eyes** — Any operation should require the confirmation of at least two people. This can help protect against ransomware and internal bad actors.
- Monitor unusual administrative activity — If an administrator is behaving outside of the norm, suspend activity until the activity can be explained.
- **Delay deletes** — In a world of deduplication, most deletion occurs only when garbage collection reaps freed blocks. Therefore, if there are excessive deletes, hold onto the blocks until the activity can be validated.
- **No root access to an underlying system** — Organizations often focus on securing the backup software management layer but forget that everything runs on a Linux or Windows box that can, itself, be compromised. If the ransomware can compromise the underlying operating system, your environment is not zero-trust.

While most organizations operationalize two or three of these requirements, the remaining exposures leave them vulnerable to attacks. As ransomware attacks become more intelligent and more aggressive, any exposure will be exploited.

Evolve

Ransomware attacks are constantly evolving. Whatever you build today will be obsolete sooner than you might expect because multiple groups are constantly releasing new ransomware packages. You need to be able to evolve with them.

Over the past few years, ransomware has evolved from:

- **Consumer to enterprise**
- **Attacking production data** to attacking backups AND production data
- **Targeting endpoints and file servers** to VMs, cloud apps, and databases

A ransomware protection solution from just two years ago is helpless in the face of a modern attack.

You are facing an army of expert attackers who spend every day trying to compromise your defenses. You can take on that fight yourself, or work with an army of expert defenders who spend every day trying to protect you. "Do it yourself" ransomware protection is not a viable option anymore. It is time to enlist a service.

> "We've slashed the time for recovery from up to eight hours down to a few minutes with Druva — about 90% faster. Our employees were used to having to wait a day for us to restore data before, and the feedback we've received since using Druva is one of astonishment at the speed of restoration."
>
> **— Tom Ferrucci, CIO at Hope Global**

druva

A Powerful Ransomware Recovery Service

Enterprises must assume that they will eventually be compromised by ransomware — if only due to human error. This means that their ransomware protection strategy must address a response plan that identifies, quarantines, and removes ransomware infections immediately, and automatically restores data to resume normal operations.

A proven vendor like Druva can help you through your business resiliency and continuity journey and provide comprehensive employee education. While no backup vendor can immunize you from future malware attacks, Druva guarantees you will significantly increase your odds of faster response and recovery. The Druva Data Resiliency Cloud empowers security operations and IT teams to protect, detect, respond, and recover faster from external or internal attacks, ransomware, as well as accidental or malicious data deletion.

Your teams will not only prevent data loss and save costs, but also accelerate response and recovery times so you can get your company back to normal in hours, not weeks or months after you've been hit by a ransomware attack.

90% faster time to restore deleted or lost files with Druva.

Prevent encryption of backups with air-gapped, immutable data and zero-trust security including MFA, SSO, and RBAC.

Contain the spread of infection APIs automate ransomware playbooks including quarantining affected resources and deleting infected snapshots.

Identify anomalous data sets and activity understand which data may be affected with unusual data activity and access monitoring.

Automate recovery of clean and complete data scan snapshots for malware *before* restoring them and automatically find the most recent clean version of each file.

druva

Next Steps

IT and infosec teams that understand the risks of ransomware are in the best position to defend their companies from attacks. By selecting the right ransomware recovery solution, you can ensure your organization has a rock-solid multi-layer defense plan in place to reduce the impact of ransomware or malware. You'll also be far less vulnerable to costly ransom demands and debilitating downtime.

Check out druva.com/solutions/ransomware/ to learn more.

druva

Sales: +1 888-248-4976 | sales@druva.com

Americas: +1 888-248-4976
Europe: +44 (0) 20-3750-9440
India: +91 (0) 20 6726-3300
Japan: +81-3-6890-8667
Singapore: +65 3158-4985
Australia: +61 1300-312-729

Druva enables cyber, data and operational resilience for every organization with the Data Resiliency Cloud, the industry's first and only at-scale SaaS solution. Customers can radically simplify data protection, streamline data governance, and gain data visibility and insights as they accelerate cloud adoption. Druva pioneered a SaaS-based approach to eliminate complex infrastructure and related management costs, and deliver data resilience via a single platform spanning multiple geographies and clouds. Druva is trusted by thousands of enterprises, including 60 of the Fortune 500 to make data more resilient and accelerate their journey to cloud. Visit druva.com and follow us on LinkedIn, Twitter, and Facebook.