TECHNICAL WHITE PAPER

Recovering Fast from Ransomware Attacks: The Magic of an Immutable Backup Architecture

Chris Wahl
September 2023
RWP-0515

Table of Contents

SUMMARY

Ransomware has been blasting my news feeds on a daily basis for years. Each article details the story of an organization that can no longer access their business critical data. Where the attackers have crippled their victims by encrypting access to production files and storage devices. According to BlackFog, a business is attacked by a cybercriminal every 11 seconds, projecting the damage costs of up to $20 billion by the end of 2021 Whilst Cyber security teams have invested in a myriad of protection tools, extortionists continue to find new mechanisms to encrypt organizations' data.

Backups are one of the most—if not the most—important defense against ransomware. But if subject to corruption, attackers will use it against you. Advanced ransomware is now targeting backups—modifying or completely wiping them out—eliminating your last line of defense and driving large ransom payouts. Rubrik's uniquely immutable file system natively prevents unauthorized access or deletion of backups, allowing IT teams to quickly restore to the most recent clean state with minimal business disruption. In fact, we are so confident in our ability to recover that we back it up with up $5M through a Ransomware Recovery Warranty. This paper walks you through our one-of-a-kind immutable architecture and robust security controls that harden your data from cyber attacks.

THE EFFECTS OF RANSOMWARE

Ransomware is designed to encrypt your data so that it is no longer usable. Often, this means encryption of data held on primary storage to overwhelm IT and require massive recovery efforts from tape or other archives. Additionally, lower level encryption of the Master Boot Record (MBR) or other operating system level encryption is used to prevent booting and other common operations. For virtualized environments, the shared data storage used to host virtual machines is a primary target, such as with NFS-backed datastores. This can effectively bring down critical services in an organization. The attackers then demand a ransom to unlock the data so that services can be resumed.

HOW RUBRIK HAS HELPED CUSTOMERS

Several customers have successfully survived a ransomware attack through the use of our immutable solution and instant recoveries as part of their defense in depth strategy.

For example, during a ransomware attack on the City of Durham, their leaders credited their quick response to Rubrik's backup solution. Durham Mayor Steve Schewell said, "The city can be assured that our backups are very good because they're immutable. [This means that] they could not be consumed by ransomware." As a result, they were able to quickly restore critical city services, including access to 911. In addition, Kerry Goode, Durham CIO, emphasized that core business systems, including ones that manage payroll, were back online by the start of the business week.

Kern Medical Center discovered a large ransomware attack had penetrated their environment when users reported they couldn't access their systems. They were able to recover 100% of the impacted systems protected by Rubrik within minutes, including recovering their business-critical electronic medical record system. CTO Craig Witmer said, "After the incident, we were so impressed that we moved more of our legacy systems to Rubrik and are fully confident that Rubrik's immutable backups will protect us from future incidents."

HOW DO YOU RECOVER FROM A RANSOMWARE ATTACK

Data backups can be an effective way to restore data that has been locked/encrypted by the attack. However, what if your backup data is also encrypted or deleted by a ransomware attack? How do you ensure that your backup data is not vulnerable to these attacks?

THE KEY IS IMMUTABLE BACKUPS

While primary storage systems need to be open and available for client systems, your backup data should be immutable. This means that once data has been written it cannot be read, modified, or deleted by clients on your network. This is the only way to ensure recovery for when production systems are compromised.

This goes well beyond simple file permissions, folder ACLs, or storage protocols. The concept of immutability needs to be baked into the backup architecture so that no security exposure can tamper with the backups.

RUBRIK IS DESIGNED FOR IMMUTABILITY

Rubrik uses an immutable architecture by combining an **immutable filesystem** with a **zero trust cluster design** in which operations can only be performed through **authenticated APIs**.

Rubrik's Immutable Architecture	Legacy/Wannabe-Immutable Solutions
Backup data is **never** exposed to external clients through insecure methods or protocols such as NFS or SMB. All operations on data have to be authenticated through appropriate credentials.	Rely on an architecture where backup software writes to backup storage through standard protocols like NFS/SMB, which use weak authentication mechanisms that can be easily bypassed.
All writes are **out-of-place**, meaning that new writes will never touch data written earlier.	Writes are done in-place, meaning there is no guarantee that the content has not changed since being ingested.
Data is fingerprinted at ingest time and the fingerprints stored along with data to ensure that once written, the data is never changed.	Fingerprints are not used to validate backup data. Restore operations blindly replace production data without verification.
Cluster communications are secured using the TLS 1.2 protocol with certificate-based mutual authentication.	Members of the cluster are given trust using a network whitelist that is vulnerable to man-in-the-middle attacks.

Rubrik's approach is in contrast with other data management systems using general purpose storage that use standard protocols such as NFS or SMB to advertise their availability to a wide assortment of clients. We often find that data management solutions using general purpose storage have limited or ineffective means for securely transacting data and, in some cases, leave files in their native format while allowing clients to read the backup data directly. This is a breach of confidentiality and puts extra burden on the customer to secure the storage independent of their data management solution.

AN IMMUTABLE DISTRIBUTED FILESYSTEM

One of our first design decisions was to construct Atlas, an immutable Filesystem in Userspace (FUSE) that was largely POSIX-compliant. This provides tight controls over which applications can exchange information, how each data exchange is transacted, and how data is arranged across physical and logical devices. Atlas is custom designed to be a distributed and immutable file system for writing and reading data for other Rubrik services.

Immutability is provided across two layers: the logical layer (Patch Files, Patch Blocks) and the physical layer (Stripes, Chunks). The dynamics between these two layers that will be explained further in the next few sections.

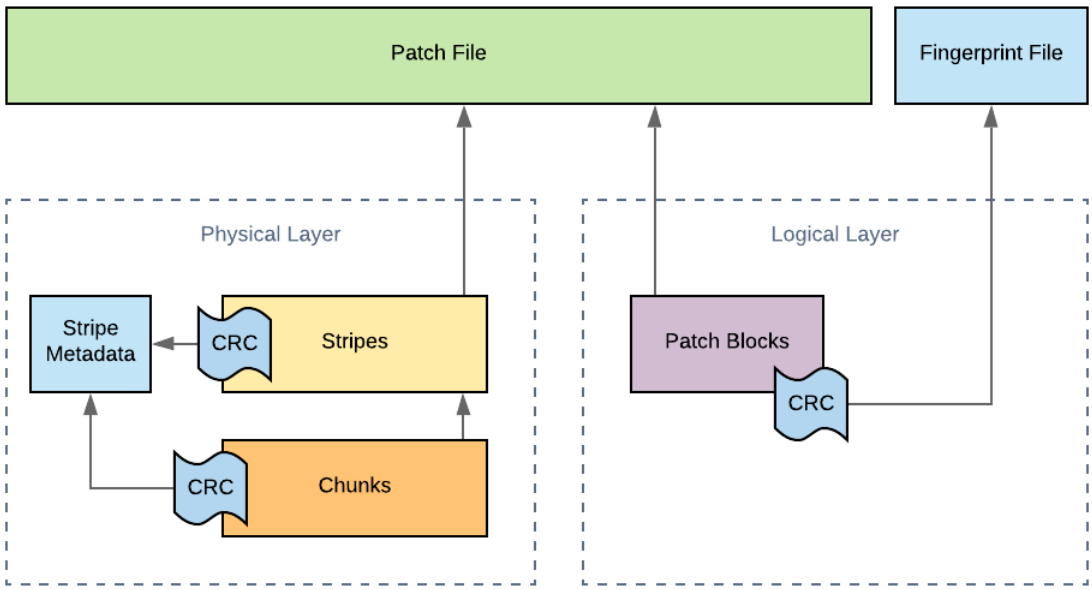

The Logical Layer

All customer data brought into the system is written into a proprietary sparse file called a Patch File. These are append-only files (AOFs), meaning that your data can only be added to the Patch File while it is marked as being open. All of the customer snapshot and journal data is held within Atlas, which enforces the use of Patch Files in the underlying directory structure. This powerful filesystem will refuse writes at the API level that are not append-only, such as situations in which the write offset value does not equal the file size. Atlas has total control over how and where customer data is written.

If your backup data has been modified, then it's essentially worthless. We solved this by ensuring that checksums are generated for each Patch Block within a Patch File. These checksums are computed and written to a Fingerprint File stored alongside the Patch File. Rubrik always does a fingerprint check before committing any data transformations. This ensures that the original file remains intact with forced validation during read operations.

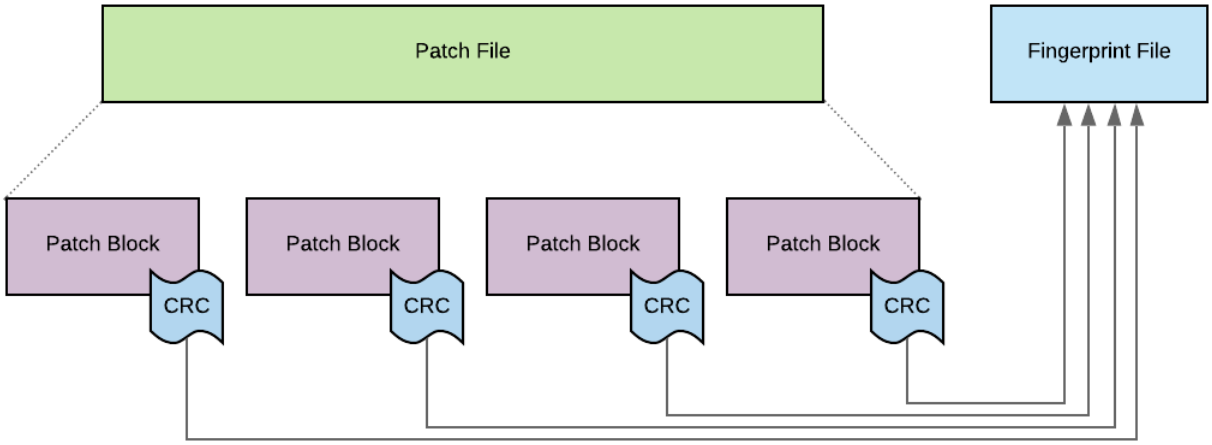

In order to counter a ransomware attack, the original, validated data must be restored from backup. Rubrik routinely verifies the Patch Blocks against their checksums to ensure data integrity at the logical Patch Block level. Patch Files are not exposed to any external systems or customer administrator accounts. This ensures that meticulous care is taken to restore exactly what you originally stored in a backup.

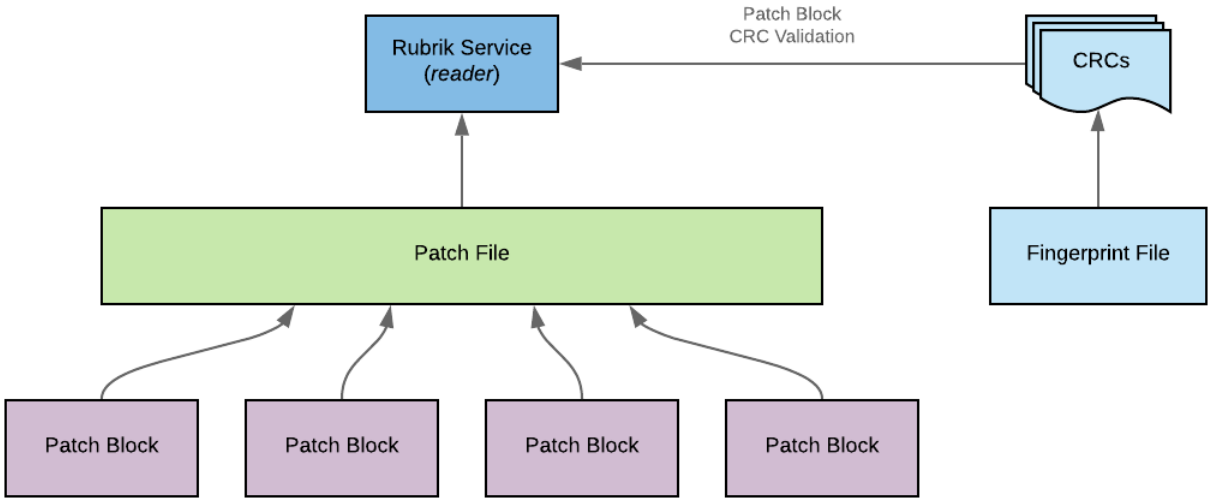

In a traditional approach, administrative access is granted to the filesystem - especially when using general purpose storage—which presents further confidentiality and integrity challenges and gives “Leakware” another attack vector. In addition, many other solutions simply restore whatever data is located in the backup folder or volume without performing validation and other due diligence on the data.

The Physical Layer

While the logical layer focuses on data integrity at the file level, the physical layer focuses on writing customer data across the immutable cluster to achieve data integrity and data resiliency. To do this, Patch Files are logically divided into fixed length segments called Stripes. As Stripes are written, the AOF computes a Stripe level checksum, which it stores within each Stripe Metadata.

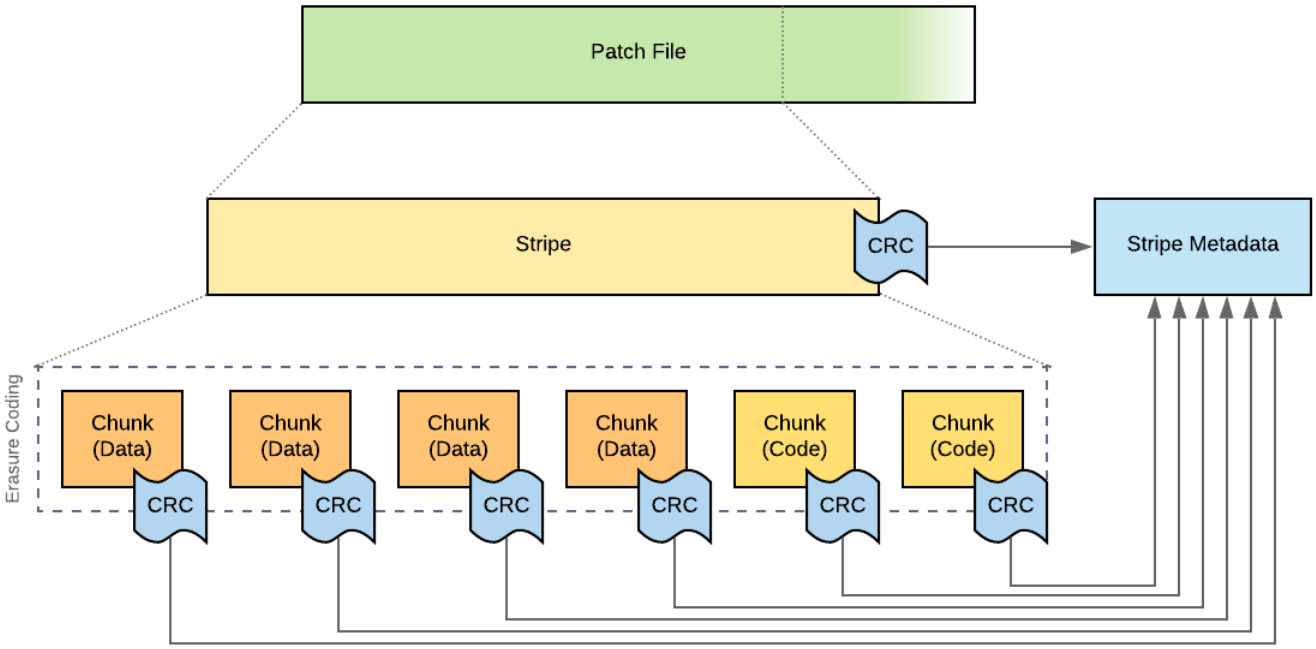

Stripes are further divided into physical Chunks stored on physical disks held within the Rubrik cluster. Activities such as replication and erasure coding occur at the Chunk level. Just as with Patch Files, as each Chunk is written a Chunk checksum is computed and stored in the Stripe Metadata alongside the list of chunks. These checksums are periodically recomputed as part of Atlas' background scan by reading the physical Chunks and comparing against the checksums in the Stripe Metadata.

Additionally, if a data rebuild is needed, the resiliency provided by erasure coding is automatically leveraged in the background.

ZERO TRUST CLUSTER DESIGN

Traditional approaches to cluster security often rely on a "full trust" model in which all members of the cluster are able to communicate with one another. In some cases, this includes root level authority, no mutual authentication checks, and the ability to read or modify your customer data that is held within the filesystem. This creates a weak surface area when designing a defense in depth architecture; if backup data can be compromised, there is no path to restoration when disruption occurs.

Secured Cluster Communications

Each cluster has some number of nodes that need to communicate with one another. This means we need to validate each node that wants to exchange data. For many solutions, there is little to nothing protecting node-to-node communication. At Rubrik, all of our intra-node and inter-cluster communication, as well as communication with external applications, use the TLS protocol with certificate-based mutual authentication for secure communication.

Rubrik does not use insecure protocols, such as NFS or SMB, to relay information within the cluster; all communication is performed through secure and trusted channels. In fact, all our internal communications use TLS 1.2 with strong cipher suites and Perfect Forward Secrecy (PFS).

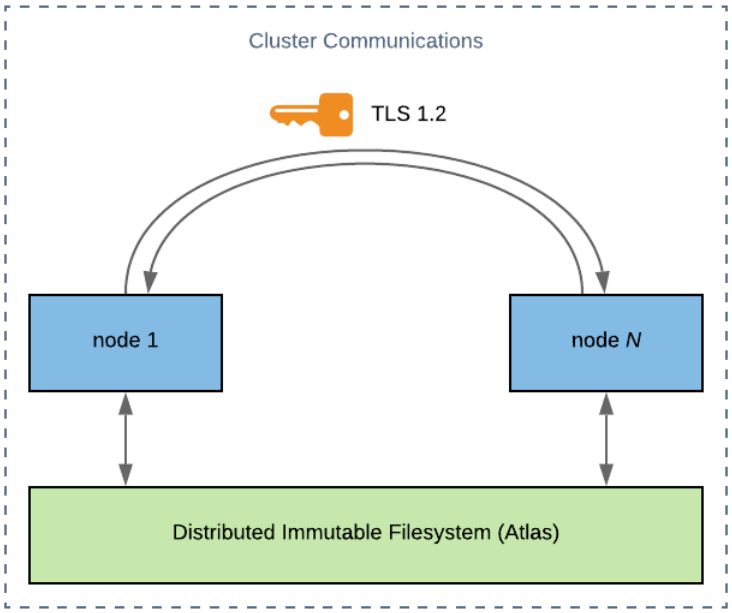

Each Rubrik cluster shipped to a customer uses strong, randomized passwords on a per-node basis. There is no concept of a "`admin/admin`" style of default local authentication that is easily searchable on the web to add an attack vector.

Systems Hardening Standards

There are numerous other elements in position to protect the integrity of the system through internal hardening standards. Here are a few that help combat ransomware:

Rubrik Hardening Standards	Key Results
All of the snapshot data held in Atlas is not exposed in a readable format via the filesystem. The only way to view the data is to properly authenticate to the Rubrik cluster, which validates that you have the correct role and permissions to view the data.	There is no way to "mount" snapshot data directly from the filesystem. Data is protected from Leakware by default.
It is not possible to run applications in a Rubrik node's kernel or user space. Only Rubrik certified services are able to run within the platform.	Solutions that allow this often provide extra attack surfaces for malicious code, human error, or other types of pain.
Rubrik pre-configures the iptables of the underlying operating system to whitelist services that can access each other.	This eliminates external access to internal services. Using a whitelist greatly reduces the attack surface area.
All Rubrik software images are signed by authorized personnel. The signature is verified during the boot process.	Trust that the software retrieved matches what was generated by the development team. Software upgrades will fail if the signature does not match.
Only the network ports that are required for user interaction with the product and communication between different internal processes are allowed. All unused ports are disabled on the product.	Ports that are not needed for the production to function are no longer potential intrusion points for attackers.

While Rubrik provides Zero Trust Data Security by default, there are a number of configurations which can be further locked down. Rubrik recommends organizations follow the constructs set forth in the Security Hardening Best Practices Guide.

AUTHENTICATED APIS

Rubrik adopted an API-first design as part of the architecture. We require authentication to all endpoints that are used to operate the solution. Authentication can be handled via credentials or secure token. This includes environments using our Role Based Access Control (RBAC) or Multi-tenancy features to logically divide the roles, features, and resources that are under management. Rubrik's CLI, SDKs, and other tools consume the API and are held to the same security requirements.

API endpoints that control the underlying behavior of the system require an additional level of authorization that can only be supplied from a certified technical support engineer. This prevents a malicious actor from being able to alter the behavior of a Rubrik cluster.

Furthermore, those APIs which can cause data removal within Rubrik are blocked by default, ensuring that even if credentials or API tokens are compromised, your data remains safe. In the event customers need to automate around the subset of blocked APIs, privileges can be granted on a token by token basis to allow scripts to consume these endpoints.

CONCLUSION

Numerous resources on the Internet advocate for a Defense in Depth strategy. This combines efforts across employee education and enablement, rapid deployment of patches, and a solid backup and recovery plan. In this post, I described how Rubrik uses a combination of data immutability and a zero trust cluster design to build a great product for protecting and recovering data. We help organizations further strengthen their ransomware response strategy through a number of value-add services. Rubrik Anomaly Detection allows organizations to quickly pinpoint exactly when attacks have taken place, along with determining the overall blast radius and scope of the attack. Rubrik Sensitive Data Monitoring then takes over, allowing organizations to quickly classify any exposed data against well known regulations such as PCI, HIPPA, or any custom, user-defined expression.

Many of our customers turn to Rubrik on their worst day. They need to be able to reliably recover from ransomware attacks to ensure minimal downtime of their critical services. A product with a truly immutable architecture provides our customers the peace of mind that when they need to, they can always access the data to recover from such debilitating attacks.

VERSION HISTORY

Version	Date	Summary of Changes
1.0	April 2020	Initial Release
1.1	November 2021	Boilerplate update
1.2	September 2023	Product naming and boilerplate updates

rubrik

Global HQ
3495 Deer Creek Road
Palo Alto, CA 94304
United States

1-844-4RUBRIK
inquiries@rubrik.com
www.rubrik.com

rwp-recovering-fast-from-ransomware-attacks / 20230920