TECHNICAL WHITE PAPER

An Introduction to Rubrik Orchestrated Application Recovery

Rubrik Technical Marketing
February 2022
RWP-0600

TABLE OF CONTENTS

RECENT HISTORY OF DISASTER RECOVERY SOLUTIONS

Ransomware and security threats continue to reach new heights annually. With a thriving underground marketplace of ransomware toolkits and "as a service" style offerings, it is no wonder why malicious software is proliferating in private networks. As IT organizations continue to work tirelessly to remediate network and software vulnerabilities, their critical data and applications are continually under attack. In parallel to these efforts, disaster recovery planning and testing have come into focus as business leadership now has a clear understanding of the ransomware risk.

Prior to the ransomware threat, disaster recovery was traditionally designed for two different scenarios: 1) application and data loss due to natural disaster or 2) significant data loss due to human error. While both of these scenarios must be addressed by a disaster recovery solution, what we will discover later is that ransomware disrupts the architectural foundations that are typically adopted by these solutions.

With data gravity at work, it is no surprise that disaster recovery solutions formed foundations around virtualization and primary storage arrays to provide application and data recoveries. This is due to the leverage that primary storage integrations provide, such as snapshotting, scheduling, and replication engines. Although this approach was effective for the aforementioned disaster scenarios of a hurricane or accidental deletion, it had many disparate components, complex compatibility matrices, and often required specialized skill sets to maintain and operate.

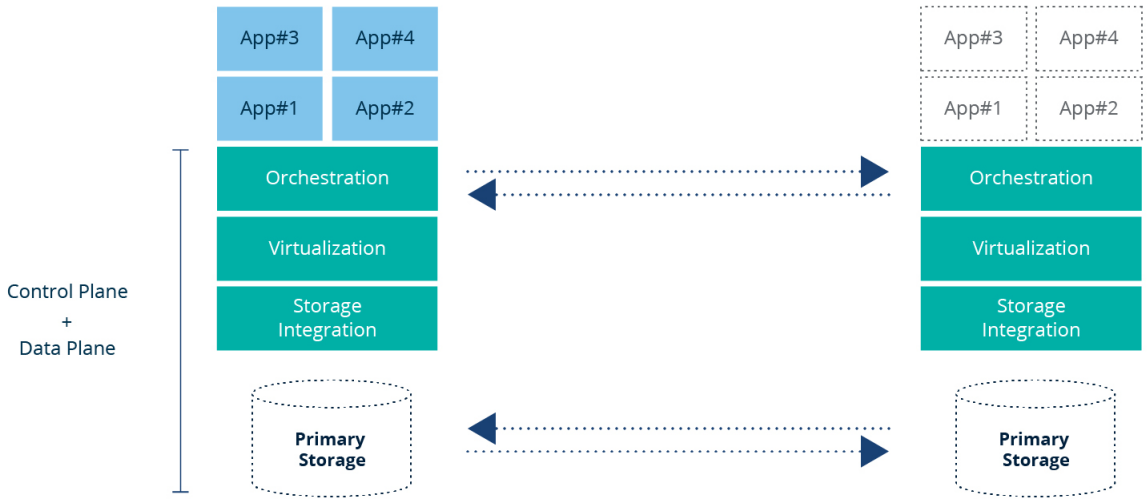

To alleviate the complexity of primary storage integrations, another architectural approach is decoupling the disaster recovery solution from primary storage. In this approach, the solution itself takes the role of the replication engine previously held by the storage provider. Additionally, it inserts itself into the data path by distributing application servers throughout the infrastructure. In doing so, the architecture adds significant technical benefits of capturing data in-flight, and journaling for granular recoveries. However, in the process of decoupling the application from primary storage, it now requires the full distribution of application servers in an environment to properly function. In essence, the complexity of integrating with primary storage is replaced with more moving parts and added complexity due to its functional requirements.

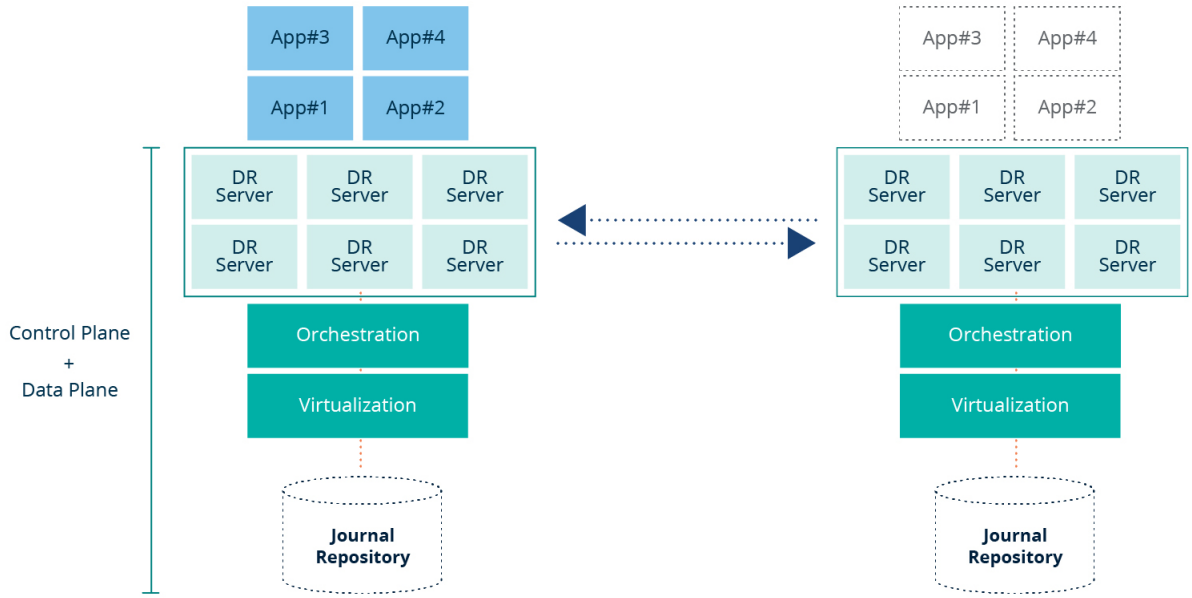

Both of these architectures are built with a dependency of either primary storage or the virtualization environment to fully operate. Unfortunately, in the age of ransomware, a third disaster scenario was introduced where a ransomware campaign can lead to the complete loss of data or data access in multiple locations, simultaneously. This disaster scenario presents significant challenges to the above-stated architectures where the ability to recover, and the time to recover, are quickly jeopardized. This is especially true if the control plane, or the recovery and orchestration plans, are in the path of the attack.

WHAT IS ORCHESTRATED APPLICATION RECOVERY?

Rubrik Orchestrated Application Recovery (OAR) is an integrated disaster recovery solution that is a part of Rubrik Zero Trust Data Security. With an intelligent design and focus on data security, Rubrik customers can simplify their disaster recovery planning, testing, and execution. In this white paper, we will discuss the Rubrik OAR design, how it works, and how it provides granular and efficient recoveries in the modern era of disaster scenarios.

HOW IT WORKS

Rubrik OAR is built and delivered as a SaaS solution to orchestrate and automate the recovery of applications running on VMware vSphere virtual machines (VM). At the fundamental level of the architecture, Rubrik OAR does not rely upon primary storage or the virtualization environment to perform orchestrated recoveries. Instead, it is a self-contained and secured architecture that has fewer dependencies that could hinder the ability to recover in the event of a widespread disaster. Added to this, the division of the data and control planes throughout the design simplifies the solution and provides added security for the configured recovery plans.

Starting with the data plane, the Rubrik CDM software and a user-defined SLA Domain Policy create a construct for data protection operations, which often include continuous data protection (CDP), backup, replication, and archival. Rubrik OAR does not initiate or influence any data protection operations, instead, it leverages them and orchestrates a specified recovery plan. In short, this simplifies the implementation and removes the need to create duplicate data streams for the sole purpose of a disaster recovery solution. Additionally, the defined granularity of recovery point objectives (RPO) from an SLA policy can be maintained without Rubrik OAR having to intervene. For example, the webserver tier may not require the same RPO granularity that the database server tier requires.

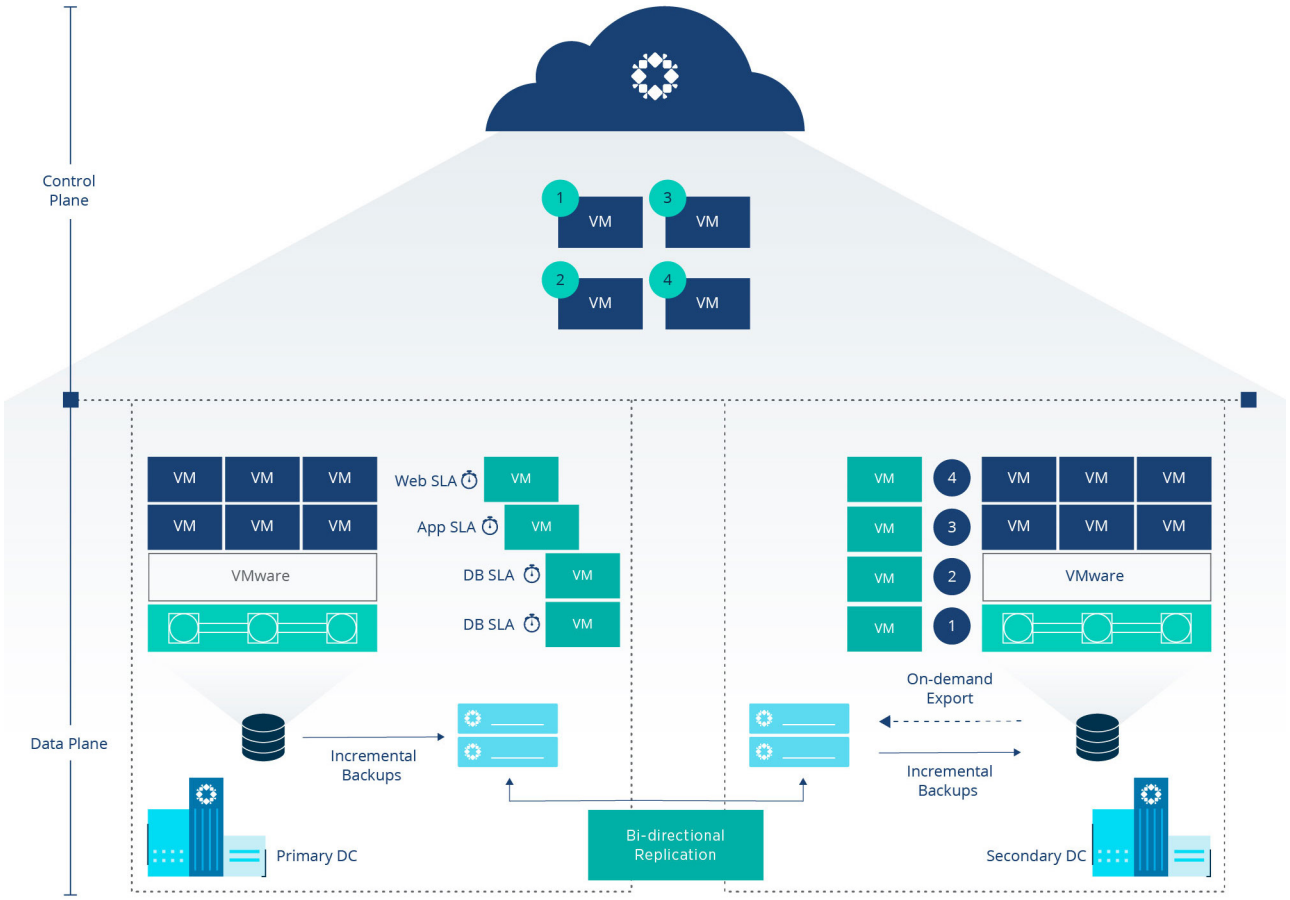

At the control plane, Rubrik OAR stores the recovery plan information and orchestrates the recoveries. By taking advantage of the interconnected architecture of Rubrik, the Rubrik OAR solution has clear visibility of the deployment and can securely broker the actions required to perform an application recovery. Furthermore, in the event of a disaster, these instructions are not dependent upon the primary storage or virtualization platforms in order to perform a recovery.

By strictly dividing the data and control planes within the Rubrik OAR solution, the approach for disaster recovery is simplified and provides a dynamic approach that can be tailored to specific requirements with minimal operational effort. Now that we have covered the high-level aspects of Rubrik OAR, let's take a closer look at the individual components that build the solution.

BLUEPRINTS

Blueprints contain the recovery plan information of all the VMs supporting an application as defined by the user. Rubrik OAR can then use a Blueprint to orchestrate either a local in-place or a failover recovery of the application. Within each Blueprint, the following items are specified from a wizard-driven interface:

- **Name –** Common name often tied to application or service.
- **Source and target CDM clusters –** The Rubrik clusters in each of the two sites are used for failover; setting the target Rubrik cluster is optional for failover enabled Blueprints.
- **Virtual Machines –** A selection of VMs protected by the same Rubrik cluster in the primary site.

- **Boot Order Priority –** Each VM is assigned to 1 of up to 5 boot order priority groups enabling orderly multi-tier application startup; more than one VM may be assigned the same priority.
- **Resource Mappings –** Compute, storage, and network mappings for each VM in the failover site.
- **RTO Optimization**
    - **Minimize Recovery Time** – Stores the latest point-in-time copy of all VM's staged (hydration; covered later) in the assigned datastore and kept up to date as new snapshots are replicated from the primary site.
    - **Minimize Datastore Usage** – Stores all the VMs on the Rubrik cluster until a recovery is launched.
- **Post-Scripts –** User-defined commands or scripts to run per VM during the recovery and configuration phase of the failover event.
- **Post Failover SLA –** Ensures backups can continue after failover; also enables replication options for further protection from site failure or to enable failing back to the primary site

Typically, each application is logically grouped by an individual Blueprint instead of the environment as a whole. This has key advantages that will be discussed later in the Rubrik Ransomware Investigation section. For now, let's dive into the mechanics of Rubrik OAR during a local, test failover, failover, and failback operations.

RECOVERY OPTIONS

After a Blueprint is defined and created in the solution, there are four primary recovery types that can be executed: local recovery, test failover, failover, and failback. All of these recovery options can be included in a single Blueprint, without having to create separate jobs, tasks, or infrastructure to execute. Let's break down each of the recovery options available.

LOCAL RECOVERY

With a local recovery operation, the solution provides the ability to recover the VMs in the blueprint to the same virtual infrastructure. The key advantage here is that the same blueprint that includes failover capabilities is leveraged for the local recovery option.

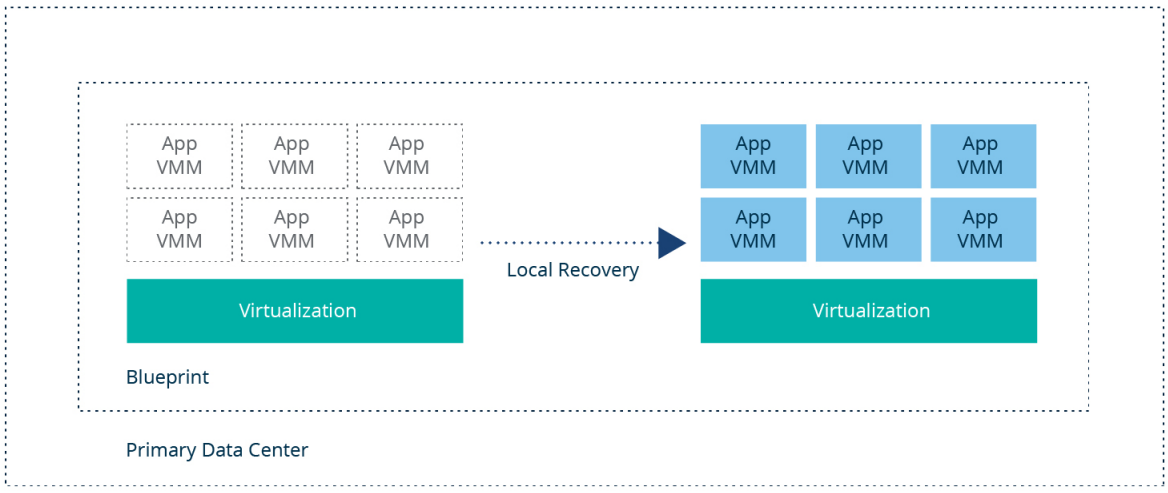

TEST FAILOVER

With the test failover operation, the VMs contained within a blueprint can be tested and validated to another data center or virtual infrastructure. It is not required to be located in a separate data center, but for illustrative purposes, it is depicted in the diagram below.

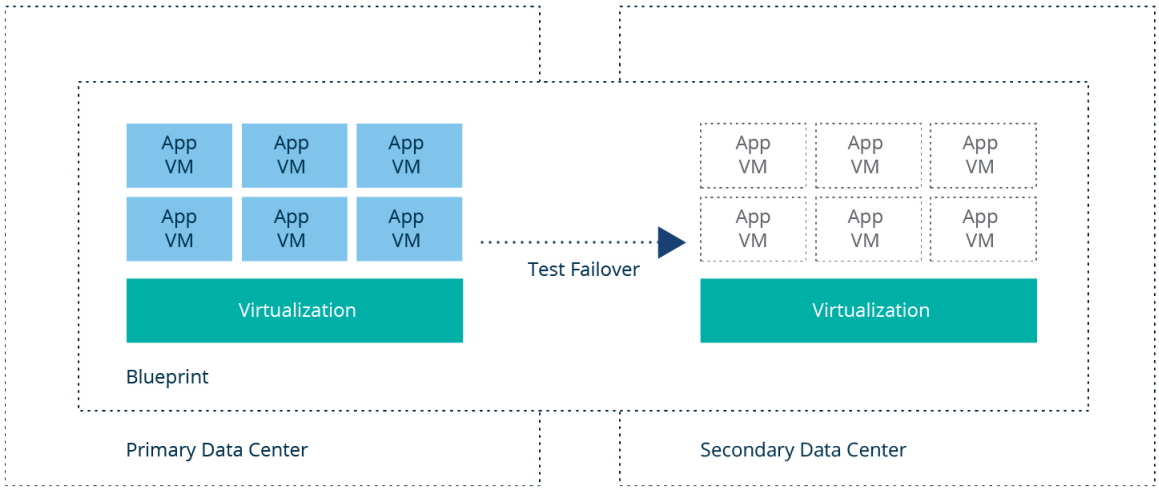

FAILOVER

When a disaster strikes, a failover operation can be executed. This operation will fully failover the application and VMs associated with the Blueprint. Additionally, the failover configurations will be appropriately applied, including IP updates and post-failover scripts as stated in the Blueprint configuration.

Continued protection for the application will be applied through the post-failover SLA that is specified in the blueprint. As a key benefit, the system does not require new full backups in order to continue protecting the application by taking advantage of VM Linking {source}. This also applies to the failback operation, which we will cover next.

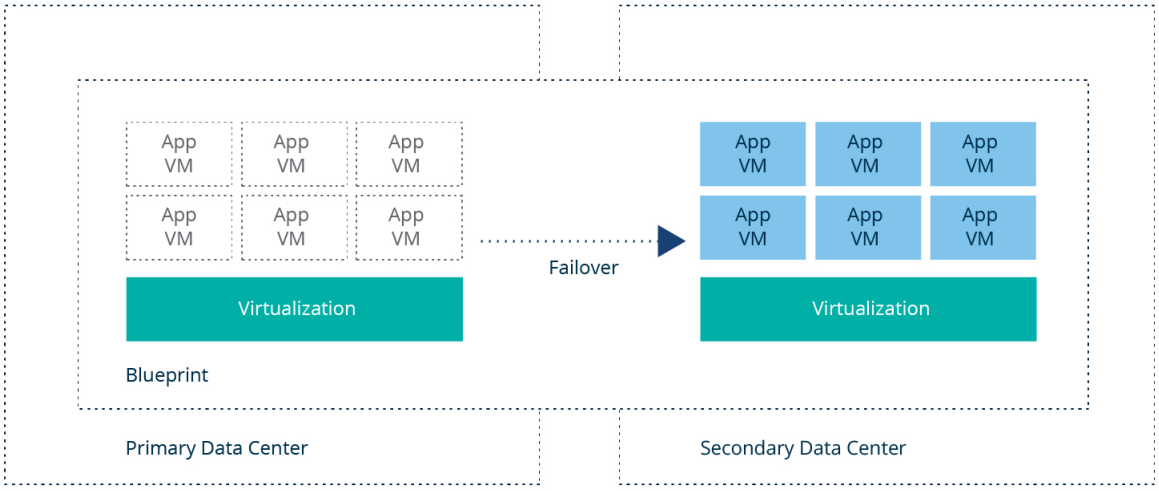

FAILBACK

In the event that the primary data center becomes readily available for use, the failback operation can be used to recover the application back to its source location. For example, if power was resumed after a lengthy outage, or an imminent threat is resolved.

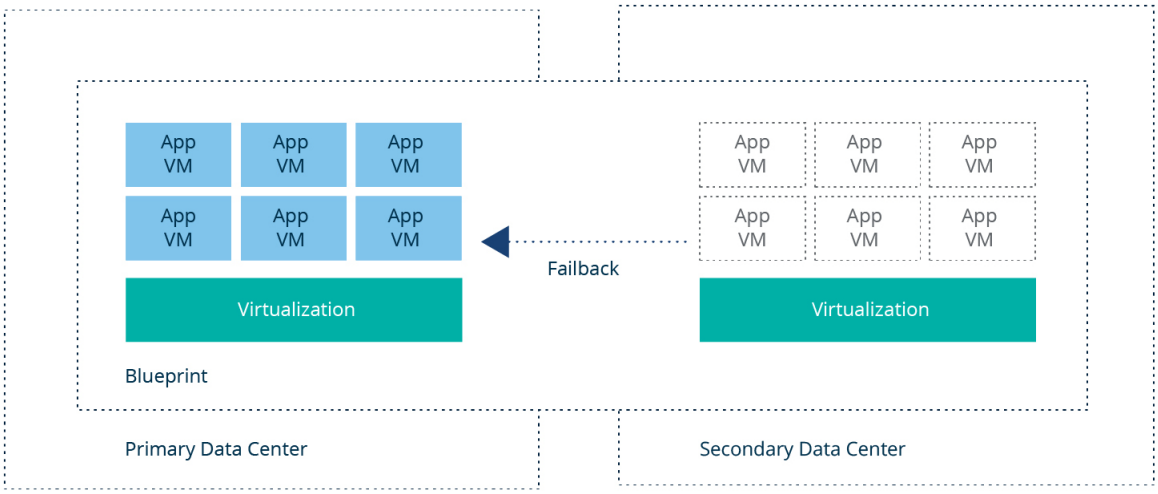

As a recap, each of these recovery options can be used through a single Blueprint where a collection of resources and configuration settings for the application are stored. This approach simplifies disaster recovery planning, testing, and execution. In the next section, we will cover the techniques and features that Rubrik uses to achieve granular recovery with low recovery point and recovery time objectives (RPO/RTO).

ACHIEVING LOW RPO/RTO

RPO/RTO are critical for application recovery, since minimizing the data loss and the amount of time an application is down, are key success factors. With Rubrik, not only is there flexibility provided by the recovery options in the Blueprint, but there are also innovative techniques that can be leveraged to achieve low RPO/RTO. Let's take a look at some of the techniques in use.

CONTINUOUS DATA PROTECTION

Although Continuous Data Protection (CDP) is a feature that has been encompassed in previous releases of Rubrik, it is a useful tool to achieve granular recovery points and low RPOs. By taking advantage of journaling, Rubrik can achieve granular point-in-time recoveries that can be replicated by Rubrik CDM and then orchestrated during recovery with OAR. Additionally, this feature can be leveraged in conjunction with regular snapshot recovery points in an application. For example, the web tier may have a static configuration where CDP does not provide an advantage. However, the application tier of servers could benefit from more granular recovery.

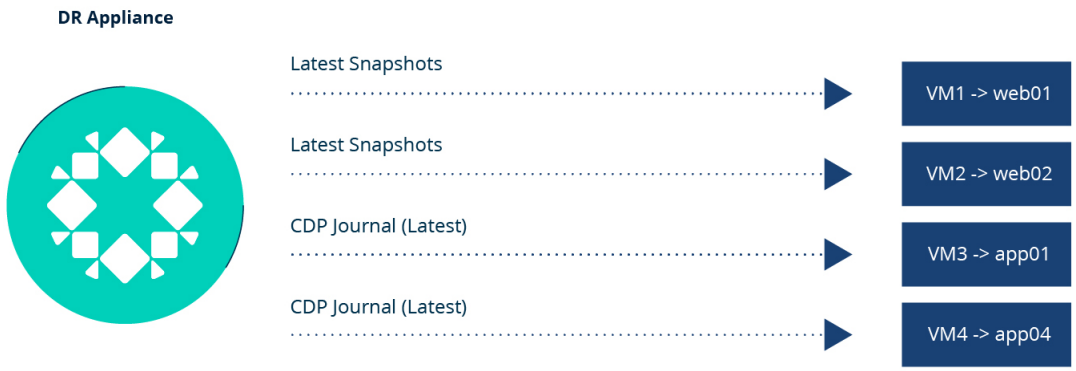

HYDRATION

Hydration is a technique used by Rubrik to pre-seed data to defined datastores that can be used for quick recoveries. In this automated operation, the backup process is essentially reversed and has two primary benefits: 1) updates are incremental to the recovery target datastore; reducing lag time for change data to be applied, and 2) lowered RTO as the target VM data is available for recovery. This operation is transparent to the end-user and created in a single Blueprint configuration where it can be used for local and failover recovery operations. A simple click of the radio button in the Blueprint will enable this feature:

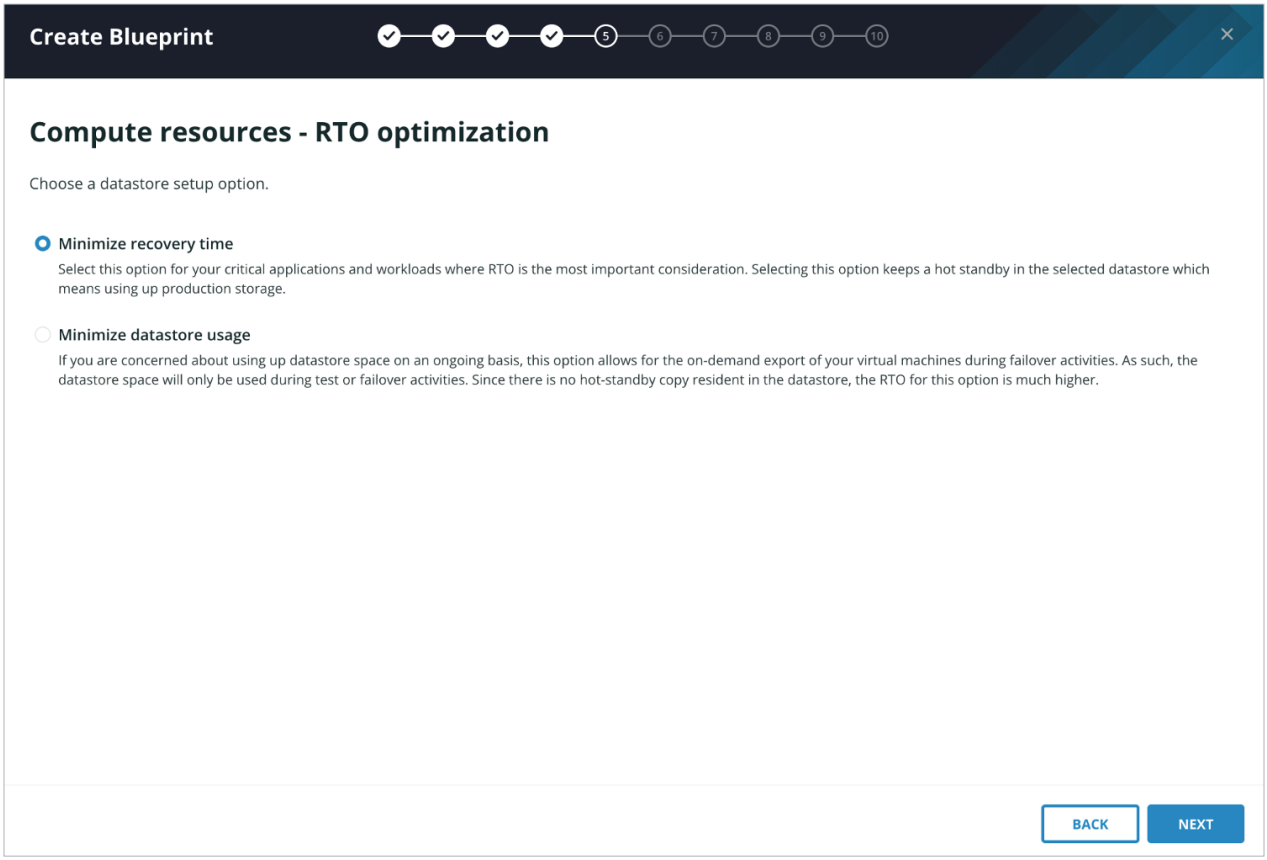

IN-PLACE RECOVERY

As previously mentioned, ransomware and cyber attacks drive the need for expedited local recoveries. With this need, Rubrik OAR can leverage a technique that can roll back a VM, or a group of VMs, to the desired point in time. This operation performs writes to the source VMs "in-place" and eliminates the need to perform an export recovery to achieve fast restoration. Let's break down the mechanics of how In-Place Recovery works, step by step.

1. User selects a previous point-in-time for in-place recovery.

2. Rubrik CDM initiates a VMware-based snapshot of the target VM.

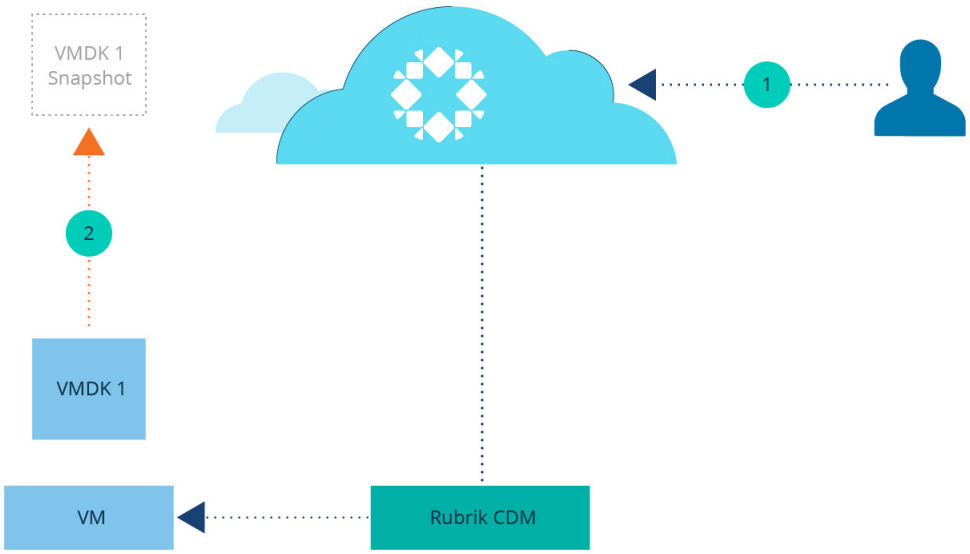

3. Snapshot 1 and Snapshot 2 are compared by Rubrik CDM and all changed blocks are identified.

   Snapshot 1: User-selected point-in-time
   Snapshot 2: VMDK1 snapshot

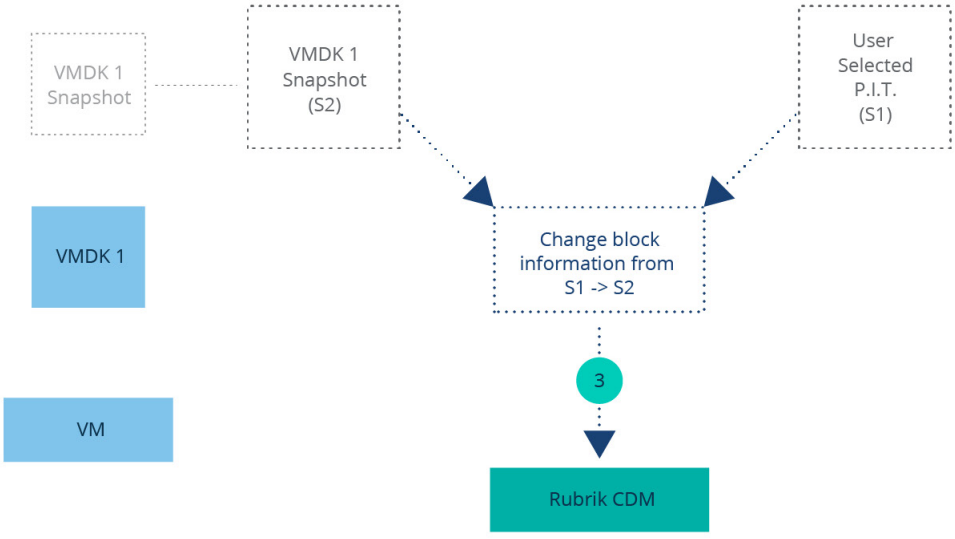

4. Rubrik CDM performs an overwrite change “in-place” of the required blocks of data to “roll back” the image to the point-in-time (Snapshot 1) selected by the user.

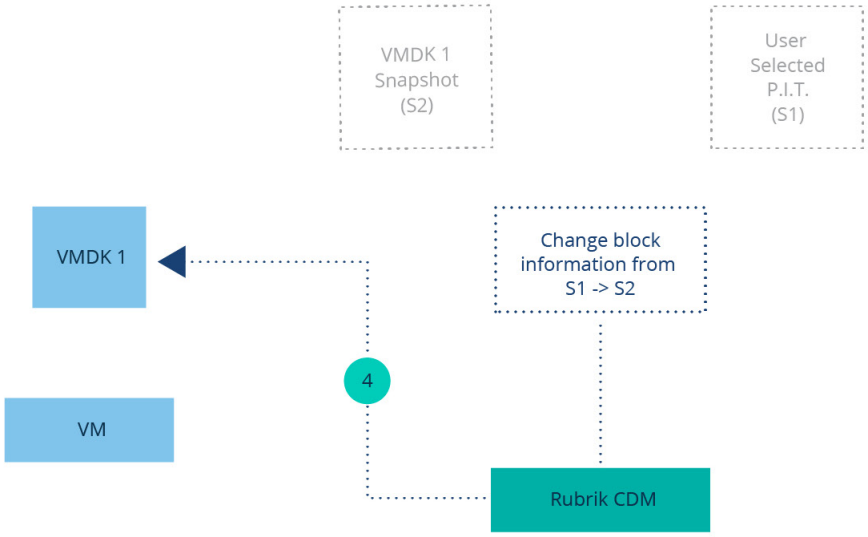

RANSOMWARE RECOVERY WITH RANSOMWARE INVESTIGATION INTEGRATION

Now that we have covered Rubrik OAR and the recovery capabilities natively in the application, let’s take a look at the integration with Rubrik Ransomware Investigation (RI) to take a closer look at how encryption and anomaly detection can be leveraged for disaster recovery during a ransomware attack. As previously mentioned, this type of disaster is relatively new to the world of IT and broadens the scope of recovery options required in a solution; not only does an organization need recovery from natural disasters with site failover capabilities, but also fast localized recoveries.

GATHER KEY INTELLIGENCE WITH RANSOMWARE INVESTIGATION

With the Rubrik OAR integration with RI, the solution is empowered to drive more intelligent recoveries with flexible options for faster recovery. In the case of a ransomware disaster, it is difficult to understand the blast radius and how it could impact your applications. By grouping your resources in a Blueprint, you can quickly identify insights that are picked up by RI.

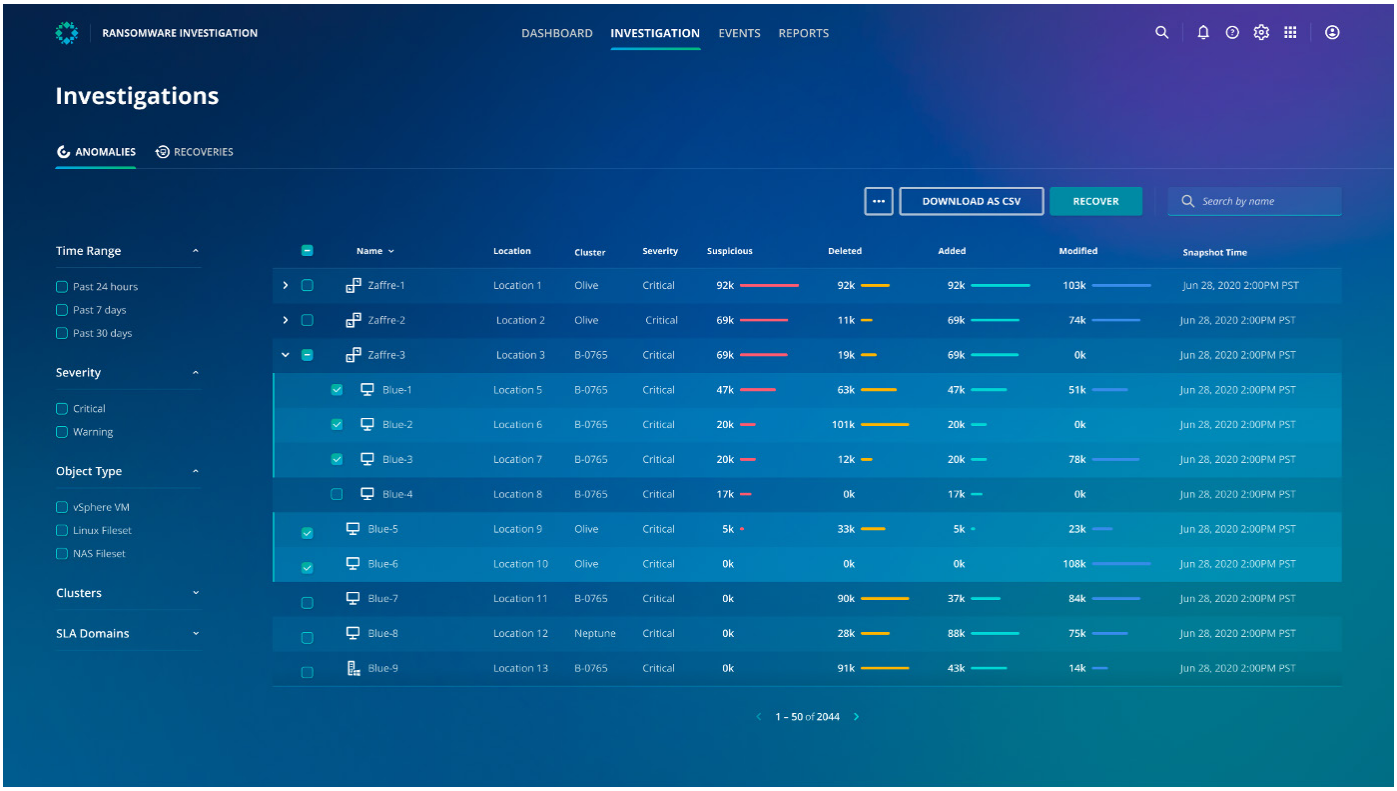

HELP IDENTIFY THE GENESIS OF THE ATTACK

Apart from the blast radius, understanding the genesis of the attack is a high priority for IT, where having flexibility in recovery options is key to expediting recovery. For example, Security Operations teams can investigate specific point-in-time copies of a system or application that can be exported from the OAR solution to a sandbox area. This is likely to be an iterative process, where multiple point-in-time copies are requested for analysis, and an opportunity to leverage the capabilities of the Rubrik solution.

HAVING OPTIONS FOR ORCHESTRATED RECOVERY

Once a point-in-time for recovery is determined to be safe to be introduced into production, the ability to use a Blueprint to expedite recoveries is key. With Rubrik OAR the customer can decide their disaster recovery path based upon their individual requirements and needs. This could include a full site recovery, or a more granular recovery of applications, in either the primary or secondary data center.

CONCLUSION

Ransomware and the modern era of IT disasters are projected to continue as a threat to the data and applications in every business vertical. It is clear that data security and being prepared are the most important IT efforts of the near-term future. With Rubrik OAR, customers can take additional steps toward a security-first approach, and be prepared to recover from disaster, when the company needs it most.

SOURCES AND NOTES

Rubrik Blog: https://www.rubrik.com/blog/products-solutions/21/5/disaster-recovery-for-vmware-applications

VERSION HISTORY

Version	Date	Summary of Changes
1.0	February 2022	Initial Release

rubrik

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

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

rwp-an-Introduction-to-rubrik-orchestrated-application-recovery / 20220622