TECHNICAL WHITE PAPER

How It Works: Rubrik Integration with Storage Arrays

Marc Creviere, Technical Product Manager
September 2023
RWP-0581

Table of Contents

CHALLENGES

Rubrik leverages built-in VMware snapshot capabilities to capture point-in-time data from your vSphere environment. When the backup is completed, the snapshot is merged and virtual disks are consolidated, resulting in what is typically an imperceptible 'stun' where I/O is paused during the final disk consolidation operation. The virtual machines are stunned for the duration of the consolidation. In typical circumstances, this process is completed almost immediately. A VM with considerable amounts of changes gathered in the temporary snapshot can be stunned for a noticeable or disruptive amount of time. This can have adverse effects on guest applications or services. Likewise, the simple existence of a snapshot can negatively impact performance of a guest, largely because of the 'write redirection' required to send new blocks to the delta disk. It is worth noting that there is a brief stun at the creation of a vSphere snapshot as well as delta disks are created, and the VM is reconfigured.

Most workloads are tolerant of this and can take advantage of Rubrik's do-no-harm approach to backup operations and native vSphere integration, but some highly sensitive applications recommend or require additional steps to minimize any impact from the vSphere snapshot process.

RUBRIK ARRAY INTEGRATION

Using Rubrik's array integration, snapshot data can be captured at the storage level and backed up by Rubrik, minimizing impact to production operations. VMs do not incur any additional latency due to the existence of snapshots while backup data is being ingested. Because the backed-up data is captured on a temporarily created volume via a proxy VM, there is no production VM stun associated with disk consolidation at the conclusion of the backup. Rubrik's forever incremental snapshot approach is still applicable after the first backup, as this integration can still make use of change block tracking. Furthermore, this integration can be enabled and disabled without impact to the forever incremental process.

In the case of Pure FlashArray, snapshot technology is non-performance impacting. Snapshots are not copy-on-write, nor are they redirect-on-write. They are simply metadata preservations of the point-in-time of the volume. No block data is moved or redirected. So there is no performance penalty when a FlashArray snapshot is created, in existence, or deleted[1] and globally data-reduced, so it does not significantly impact operation or capacity during backup operations. For information about other supported vendors array snapshot and performance related details, please refer to your array vendor's product documentation.

Rubrik's array integration is an effective way to protect workloads that are specifically vulnerable to latency caused by vSphere snapshot stun time. For all other workloads, Rubrik recommends that standard protection workflows be used.

ARCHITECTURE

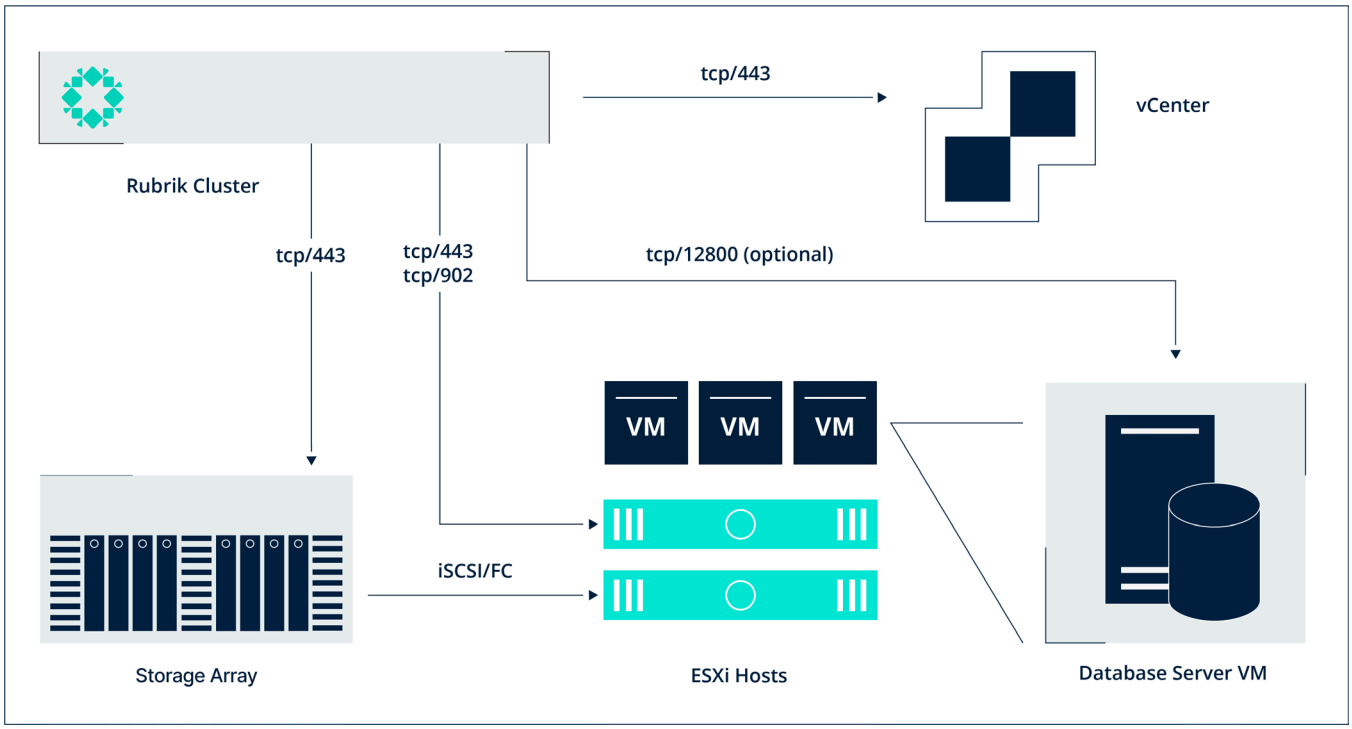

Rubrik will communicate with vCenter and the storage array using their respective APIs over HTTPS. Snapshot data will be captured from the ESXi host via NBDSSL, while pre and post scripts will either run via RBS directly to the guest OS or via VMware Tools accessed through the hypervisor host[2]. Rubrik does not place initiators/targets directly in an iSCSI or FC fabric.

PROTECTION WORKFLOW

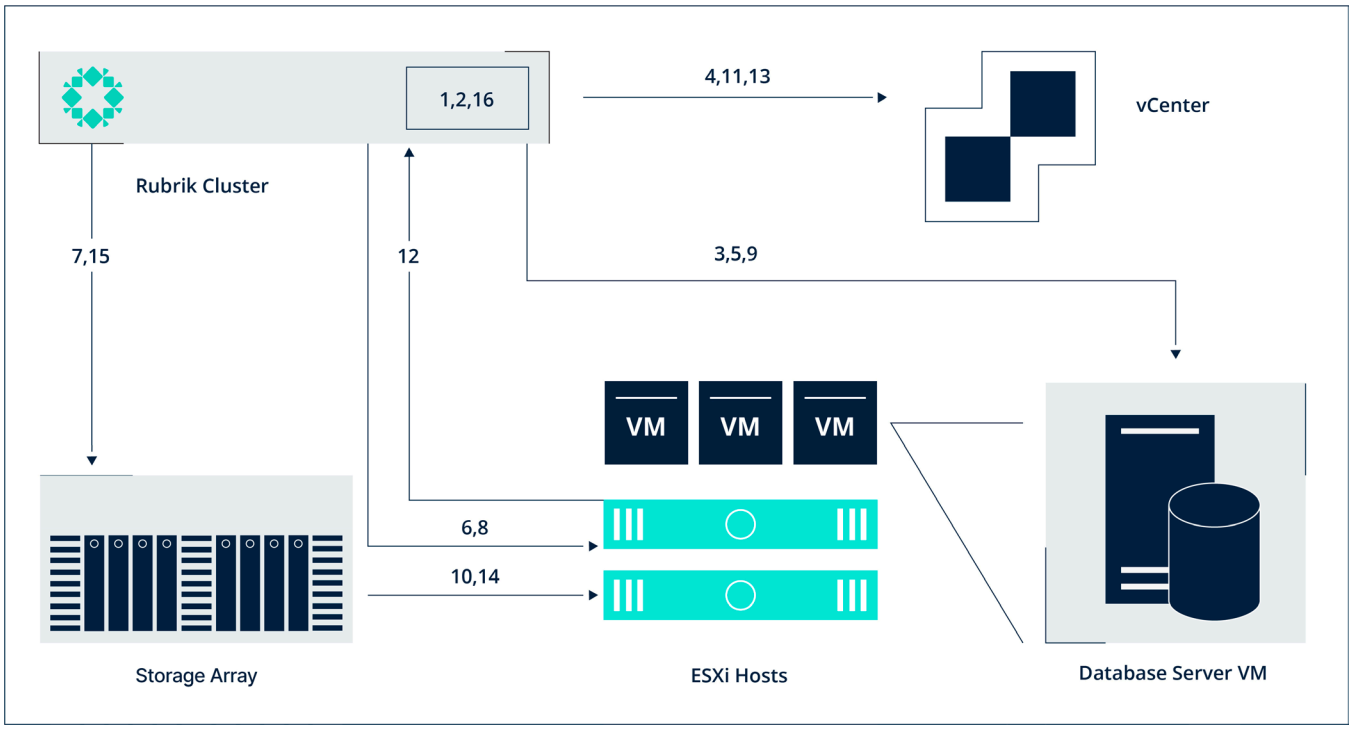

1. Backup is scheduled as per SLA or manual snapshot request.
2. Rubrik node is assigned to the backup job.
3. Guest authentication is verified.
4. Rubrik queries vCenter to identify the vSphere host containing the backup source.
5. (If configured) Pre-script is executed (database 'freeze' as example)
    - If Rubrik Backup Service (RBS) is installed on the guest operating system and connected to Rubrik, this will be used to call scripts within the guest. RBS runs as a service within the guest operating system. Communication between Rubrik is secure and does not require credentials stored in Rubrik to be used at the time of script execution.
    - If RBS is not installed, then VMware tools via the ESXi host will be the path for script execution within the guest, and will be authenticated at the time of execution using guest credentials stored on the Rubrik Cluster.
6. vSphere snapshot is created.
    - If RBS is installed, then it will act as the VSS provider, falling back to vSphere tools as the VSS provider if RBS is not installed. Using one of these tools the operating system is quiesced, or prepared for backup by flushing unwritten data from memory to disk or any application-specific tasks required to ensure that data on disk is in a consistent state for backup. For more information about Rubrik RBS, see Rubrik's blog titled Why We Built Our Own VSS Provider.

7. Array creates volume snapshots for any volume containing data for the source VM.

    - In the event that Rubrik is unable to contact the array or unable to create and mount a volume snapshot, the default behavior is to continue the backup without array integration.

8. vSphere snapshot deletion is initiated and virtual disks consolidated.

    - Because the snapshot only existed for a very short amount of time and I/O has optionally been frozen via pre-execution script, stun time and subsequent impact to the running workload is minimized.

9. (If configured) Post-snap script is executed (database 'thaw' as example).

    - Steps 6-8 take place in a matter of seconds, resulting in minimal snapshot delta to be merged and leaving production I/O free to resume while subsequent data protection takes place via a proxy VM (see below) and without interaction with the production guest VM. It is important to note that while snapshot deletion is initiated within seconds, full consolidation is dependent on change data created during steps 6-8. Pre- and post- script duration are dependent on script execution time. **This is where the reduced impact to production vs. normal snapshot operation is fully realized.**

10. Array snapshot is mapped to vSphere cluster hosts and mounted as vSphere datastore(s) if data spans multiple volumes.

    - Cluster storage is re-scanned during datastore create/remove operations.

11. A proxy VM with random naming prefix is created and virtual disks from snapshot volume(s) are attached.

    - The proxy VM is a temporary shell VM that is created with the express purpose of attaching the virtual disks from the snapshotted volume(s) via VMware vStorage APIs for Data Protection (VADP). It remains powered off and thus consumes no CPU or memory resources on the ESXi host/cluster.

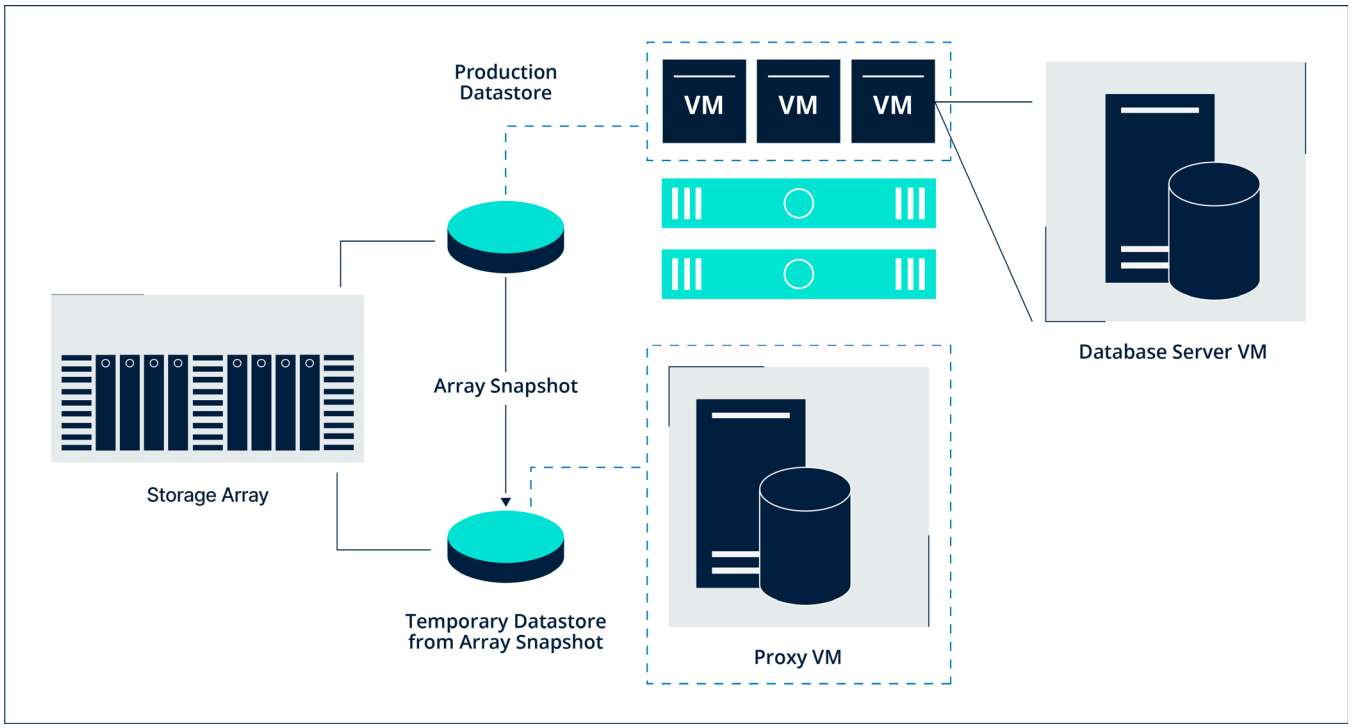

12. Rubrik ingests backup data via proxy VM.
    - Since the snapshot was taken on the production VM, CBT (change block tracking) is preserved and all backups after the first are appropriately incremental.
13. Once backup data is collected, proxy VM is deleted from vSphere.
    - Because a proxy VM is being used, there is no VM stun due to disk consolidation affecting the production workload at completion of backup.
14. Array snapshot volume datastore is unmounted from vSphere.
15. Array snapshot volume is removed from the array.
16. Backup workflow is completed.

SNAPSHOT STORAGE AND RECOVERY

Once data has been protected with this method, it functions as any other Rubrik-protected virtual machine snapshot. Data will be indexed and can be replicated, archived, and restored using any of Rubrik's standard processes.

LIMITATIONS

- Continuous Data Protection (CDP) is not compatible with this workflow.
- Guests stored on Virtual Volumes (vVols) are not compatible with this workflow.

Note: VMware vVols utilize array snapshots in the background for snapshot operations. Virtual machines backed by vVols datastores still experience stun at snapshot creation and deletion, but are unaffected by performance impact due to the existence of a snapshot.

CONFIGURATION

SLA DOMAIN

This workflow does not require any changes to the SLA domain, except for requiring that CDP is not in use. Workloads protected with this workflow can be archived and replicated as normal.

PRE/POST SCRIPTS

Any Pre-and-post scripts recommended by the application vendor, such as Caché database freeze/thaw scripts, should be configured on the guest prior to backup operations. In order to ensure proper authentication to run scripts, either the Rubrik Backup Service (RBS) should be installed within the guest operating system, or guest authentication configured within Rubrik. When using RBS, the script will be executed using the permissions of the service account configured for RBS.

WORKLOAD PLACEMENT

Virtual Machines to be backed up using this workflow must be configured with all disks located within datastores on the same supported array. Refer to the compatibility matrix on the Rubrik Support Portal for the most up-to-date information on supported arrays.

PERMISSIONS

VMware vCenter must be configured using an account that meets the minimum requirements for Rubrik as specified in the Rubrik user guide.

Refer to the Rubrik User guide for any guidance on accounts related to specific vendor arrays.

ENABLING ARRAY INTEGRATION

Array Integration is a per-VM setting and must be enabled for each virtual machine requiring this workflow. Once enabled, the array integration workflow will be used for all subsequent snapshot operations.

CONCLUSION

By utilizing Rubrik's array integration, specific workloads with especially high sensitivity to latency can be protected without the impact that can result from the existence or removal of a vSphere snapshot. By leveraging Rubrik-orchestrated guest, host, and array operations, the data capture will be offloaded to a proxy VM and the production server can continue to run without performance impact. This allows applications such as the Caché database used by Epic Systems Electronic Health Record system to be backed up as a full VM in a way that ensures maximum performance and availability.

While this functionality serves to reduce impact to the virtual machine it is configured for, it is **not recommended as a default methodology for protection at scale**. Array integrated snapshots should be used selectively where specifically required, as it does nominally increase operational overhead of protecting the vSphere environment.

Additionally, VMware made improvements to VMFS snapshot consolidation in VMFS 6, so if the environment being protected still has VMFS 5 datastores efficiencies can be realized for both backup methodologies by migrating to a VMFS 6 datastore.

SOURCES AND NOTES

1 Rubrik and FlashArray Integration: Why it matters - Cody Hosterman, Pure Storage
https://blog.purestorage.com/purely-technical/rubrik-and-flasharray-integration-why-it-matters/

2 Overview of VSS, Application or Filesystem Consistent Snapshots.
https://support.rubrik.com/s/article/000001619

ADDITIONAL INFORMATION

Overview of virtual machine snapshots in vSphere
https://kb.vmware.com/s/article/1015180

Best practices for using VMware snapshots in the vSphere environment
https://kb.vmware.com/s/article/1025279

VERSION HISTORY

Version	Date	Summary of Changes
1.0	May 2021	Initial Release
1.1	February 2023	Updated to accommodate additional array vendors and RSC.
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-hiw-rubrik-integration-with-storage-arrays / 20230913