White Paper

Why NetApp ONTAP for Kubernetes?

Justin Parisi, NetApp
January 2023 | WP-7360

Abstract

This document covers why Kubernetes and containers run seamlessly and efficiently, as well as the numerous benefits of using NetApp® ONTAP®.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Why NetApp ONTAP for Kubernetes?

NetApp® ONTAP® is a versatile, flexible, robust storage operating system that delivers high availability and enterprise-class data protection functionality to critical data storage sets, regardless of the application running in the background.

That includes Kubernetes and containers. In fact, Gartner rates NetApp's AFF A-Series running ONTAP the highest of all vendors in service scores for containers, and rates Cloud Volumes ONTAP in the top 15.

Figure 1) Gartner service scores for containers.

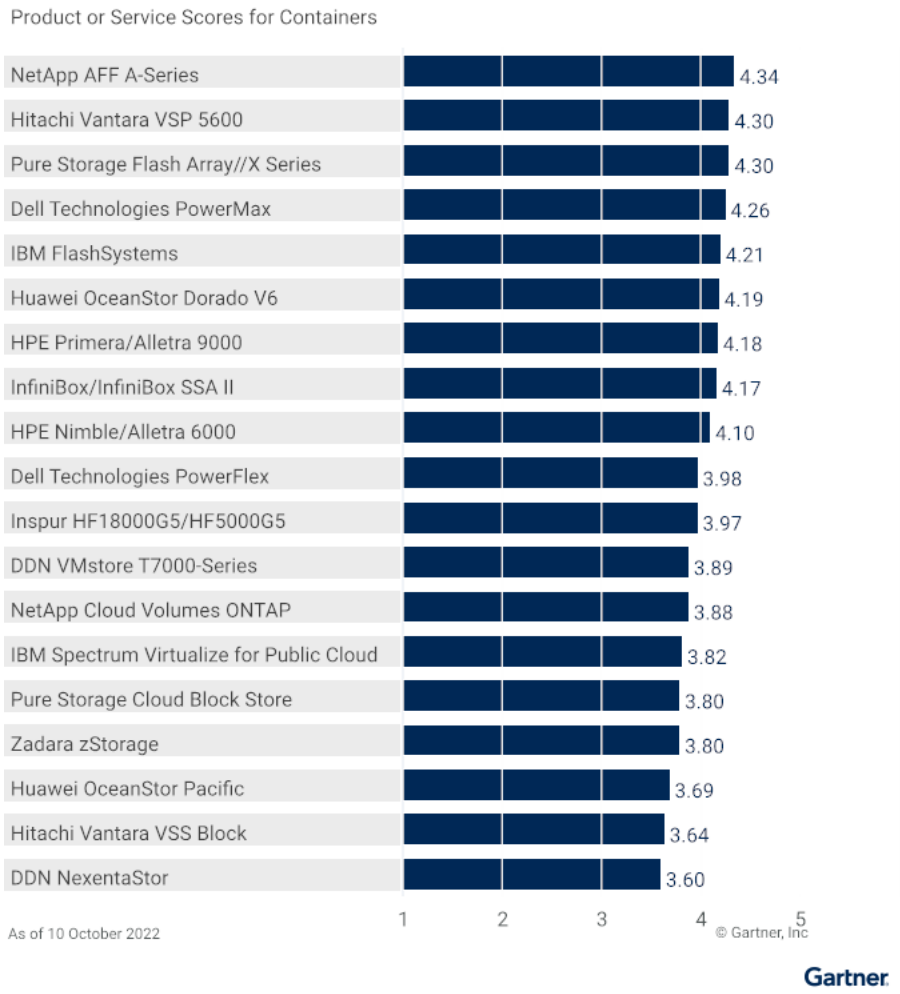

Gartner's critical capabilities are defined on their site and include the following:

- **Performance.** IOPS, throughput, latency, performance scale.
- **Storage efficiency.** Deduplication, compression, auto-tiering.
- **RAS (reliability, availability, serviceability).** Failovers, replication, hardware error checking.
- **Scalability.** Ability to grow or shrink capacity and performance as needed.
- **Ecosystem.** Support for a wide range of applications, operating systems, and cloud vendors.
- **Multitenancy and security.** Isolation of workloads, user access controls, logging.
- **Manageability.** Single management panes, automation, monitoring.

A score is derived based on how well the storage solution addresses the capabilities, and then categories are weighted based on each capability. Figure 2 shows how Gartner weighs each capability as they pertain to container workloads.

Figure 2) Gartner critical capability score weights for containers.

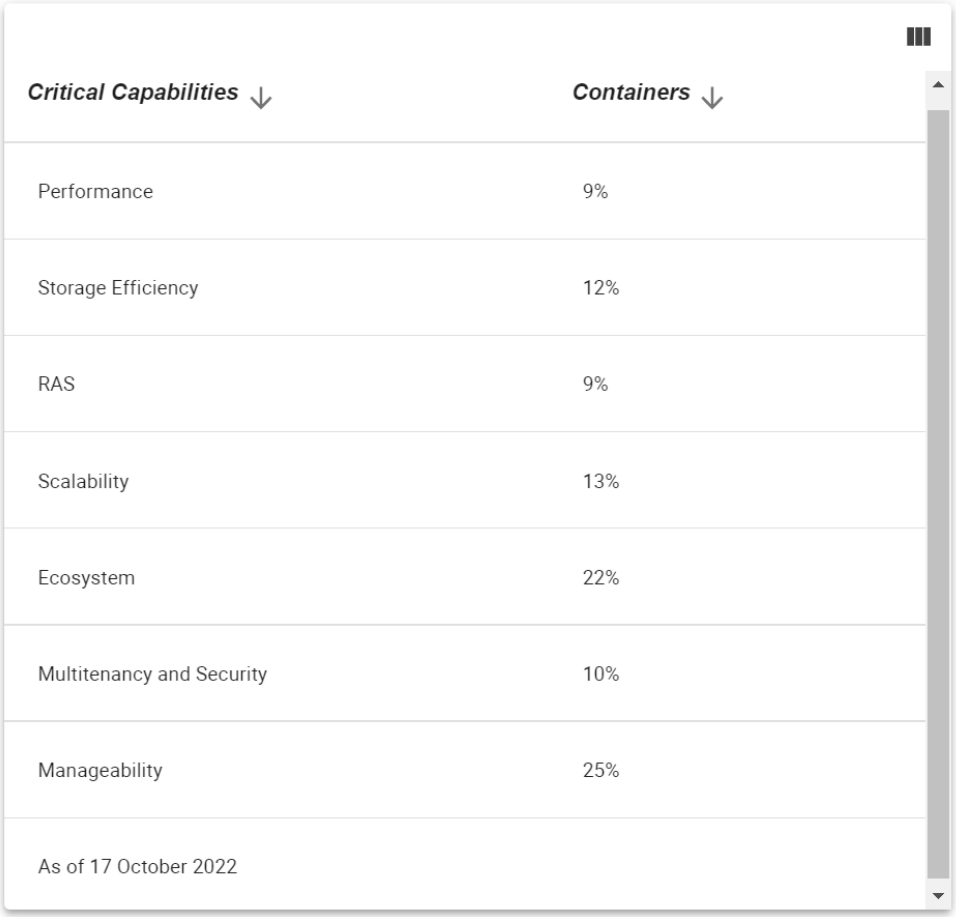

Critical Capabilities	Containers
Performance	9%
Storage Efficiency	12%
RAS	9%
Scalability	13%
Ecosystem	22%
Multitenancy and Security	10%
Manageability	25%
As of 17 October 2022	

Source: Gartner, Inc.

In Figure 2, for instance, performance carries just a 9% weight for performance, while things like ecosystem and manageability are much more important factors in the overall scores (1-5, where 5 is the best).

Table 1 shows a side by side of Gartner's critical capabilities, the NetApp score, and some ways in which NetApp ONTAP fulfills those needs.

Table 1) Gartner critical capabilities and how NetApp ONTAP fulfills those needs.

Gartner capability	NetApp score	NetApp ONTAP features
Performance	4.1	• All-Flash FAS • FlexCache® and FlexGroup • NVMe over Fabrics
Storage efficiency	4.3	• Deduplication, compression • Temperature-sensitive storage efficiency • FabricPool

Gartner capability	NetApp score	NetApp ONTAP features
		• FlexClone®
RAS	3.9	• High-availability cluster pairs for nondisruptive failovers • Snapshot™ and SnapMirror® • MetroCluster • SnapMirror Business Continuity
Scalability	4.2	• Adaptive QoS • FlexGroup volumes • Volume moves
Ecosystem	4.6	• Multicloud approach to storage • Application-aware backups • Strategic partnerships with major server and application vendors
Multitenancy and security	4.2	• Storage virtual machine architecture • Easy FIPS compliance • AES encryption for in-flight and at-rest data • FPolicy® and native file auditing
Manageability	4.5	• AIOps approach to storage monitoring with Active IQ® and ONTAP Cloud Essentials • REST API support for easy integration into standard automation such as Ansible • NetApp Astra™ Control for simplified Kubernetes integration

How ONTAP can address Kubernetes workload requirements

Although Kubernetes datasets have traditionally been stateless and containers are ephemeral, more and more applications are starting to use stateful data for their Kubernetes-based workflows. As a result, those datasets need to reside on secure, resilient storage systems that can be provisioned both on premises and in the cloud, while providing ways to back up and protect persistent volume claims (PVCs) for Kubernetes clusters.

Table 2) Kubernetes workload requirements with corresponding ONTAP features.

Kubernetes workload requirement	ONTAP features to use
Customized storage classes	• Storage virtual machine level QoS and quota settings. • Qtree QoS and quotas.
Multiple pods writing to the same PVCs	• NFS exports • Multiprotocol NAS (NFS and SMB) • Multiprotocol S3 (NAS shares presented as S3 buckets)
Support for different workload access controls	• Granular ACLs • Granular export policies and rules • Ability to specify different mount options for PVCs
Data protection for Kubernetes stateful data	• SnapMirror protection with Astra Control support
Dev/test workloads using CSI volume clones	• FlexClone
Capacity savings for duplicated files	• Deduplication • Compression • FabricPool (cloud tiering)
Scalable storage for cloud-native K8s engines for data locality	• Cloud Volumes ONTAP service (AWS and GCP)

Kubernetes workload requirement	ONTAP features to use
	• Azure NetApp Files • Cloud Volumes ONTAP (self-managed) • Amazon FSx (self-managed) • NetApp FlexCache volumes
Apps with different access mode requirements	• iSCSI (RWO) • NFS (RWX)

Why external storage instead of local disk for Kubernetes deployments?

Many Kubernetes cluster deployments use local disk on the worker nodes for storage. In many cases, this works fine, but when a cluster or application starts to scale, the local disk can become a bottleneck. In addition, if you need multiple clients to be able to access that data over a network at the same time, then issues such as disk utilization, CPU utilization on the worker node, and lock contention will create potential bottlenecks or even data consistency issues. Local storage also doesn't leverage concepts such as point-in-time snapshots or native backup capabilities, so data management becomes a headache when dealing with mission-critical datasets.

Scale, data protection, and performance are just a few of the reasons to look at an external storage provider for Kubernetes datasets. Table 3 shows a side-by-side comparison of a handful of challenges that NetApp ONTAP can address when compared to local disk.

Table 3) External storage with NetApp ONTAP versus local disk for Kubernetes clusters.

	Local disk	NetApp ONTAP
Ability to scale with applications	N	Y (up to 24 nodes)
Independently scale compute and storage	N	Y
Cross-region or cross-data-center DR (cloud and on premises)	N	Y
Reduction of CPU and disk bottlenecks	N	Y
Ability to leverage hybrid disk configuration (SSD and spinning disk)	N	Y
Multiple protocols for client data access to the same datasets	N	Y
Ability to handle massive amounts of simultaneous client access	N	Y
Hybrid cloud functionality	N	Y
Storage efficiencies (deduplication, cloud tiering)	N	Y
Space-efficient clones and Snapshot copies	N	Y
Nondisruptive operations (upgrades, maintenance)	N	Y
At-rest granular data encryption	N	Y
QoS controls	N	Y

Flexibility of deployments

Modern Kubernetes application deployments can reside in the cloud, on premises, or on both in a hybrid approach. Whether you're using a Kubernetes engine in the cloud, such as Google Kubernetes Engine (GKE), Azure Kubernetes Service (AKS), or Amazon Elastic Kubernetes Service (EKS), or rolling your own deployments on premises, it's important to have your datasets as close to the K8s cluster as possible to improve performance and maintain reliable connectivity. ONTAP backed PVCs can be deployed in on-premises data centers or in any cloud that uses managed services like NetApp Cloud Volumes Service for Google and AWS, Azure NetApp Files, or self-managed deployments using Cloud Volumes ONTAP, or Amazon FSx for NetApp ONTAP, so regardless of where your Kubernetes cluster resides, you'll have ONTAP right by its side.

If you can't stand up an entire ONTAP cluster local to your Kubernetes deployments, you can deploy sparse caches using FlexCache for NAS access in multiple locations for fast, localized reads of large datasets. With ONTAP, your Kubernetes clusters have many options for stateful data access.

Reliability and resilience

When running your business-critical applications in Kubernetes, the last thing you want to worry about is the availability of your data. With ONTAP, data access is reliable, with "six-nines plus" availability (as per this IDC white paper). That is achieved through a combination of features and functionality, such as:

- Redundancy in power and storage pathing
- RAID protection from multiple drive failures
- Hot-swappable components
- Checksum error correction
- Nondisruptive upgrades
- Rapid data cloning
- Snapshot copies
- SnapMirror replication for DR and backup
- High-availability pairs for storage failovers
- Cyber-resilience via autonomous ransomware detection*
- Integration with NetApp Astra for full Kubernetes cluster backup and restore
- Space- and network-efficient incremental forever backup to object stores using SnapMirror Cloud.

* Currently not supported with NetApp Astra Trident.

Not only is data protected from failures with ONTAP, it is resilient in the rare instance that an outage does occur. Hardware failures result in nondisruptive storage failovers. Deleted or corrupted files – or even ransomware-locked files – can be restored in seconds using Snapshot copies. Entire Kubernetes cluster namespaces can be failed over to DR sites in the event of a total loss of the primary data center by using a combination of ONTAP SnapMirror data protection and Astra Control for Kubernetes backups.

With ONTAP, your Kubernetes applications can withstand many of the worst-case scenarios that bring businesses to their knees and can reduce the risk of revenue loss in those situations.

The following sections cover the unique benefits of both cloud and on-premises deployments of ONTAP for Kubernetes.

Robust automation and analytics support

When hosting data for a Kubernetes workload (or any workload, really), easy-to-consume visibility into critical storage statistics and the ability to automate repeatable tasks are important to any storage, application, or DevOps engineer. With ONTAP, there are many ways to manage your cluster, through a simple GUI or through standard REST API support.

- For on-premises storage, ONTAP System Manager offers not just a way to provision and view storage, but also real-time performance statistics that are kept for up to 1 year for historical trending. File Systems Analytics show file-level information such as size, location, access/modification times, and directory sizes, and whether or not your volumes are protected in case of a disaster. System Manager supports full REST API integration for the automation suite of your choice and contains comprehensive documentation around the supported REST API functions available.
- For cloud instances of ONTAP, NetApp BlueXP™ provides IT experts and cloud architects with a centralized control plane to manage, monitor, and automate data in hybrid multicloud environments.
- In addition to ONTAP System Manager, NetApp offers Active IQ Unified Manager and Active IQ Digital Advisor, which deliver even deeper levels of performance information with longer retention

times and best practice recommendations that include “Fix It” buttons for fast correction of common storage issues. While ONTAP System Manager is dedicated to a single cluster, Active IQ can support multiple ONTAP clusters, both on premises and in the cloud.

- NetApp Cloud Insights extends what Active IQ can do and presents it in an “as a service” format so that you don’t have to install or manage any server hardware or VMs. Instead, Cloud Insights is cloud native and is automatically upgraded and maintained while offering ransomware detection*, ONTAP Essentials, and integration for Kubernetes clusters and Astra Control for a more universal approach to managing DevOps environments. For more information on ONTAP Essentials, check out Episode 344 of the Tech ONTAP Podcast.
- NetApp also offers a robust set of Ansible modules for common storage admin tasks so you don’t have to reinvent the wheel when setting up your automation suite to include ONTAP storage.

* Currently not supported with NetApp Astra Trident.

ONTAP for Kubernetes benefits: On premises

When running ONTAP in your on-premises data centers, you get the full suite of ONTAP features and functionality, as well as complete control over the storage infrastructure. Although that means more overall management is needed, it also means that you get to choose the platforms, configurations, and automation and customize them for your specific needs. For Kubernetes and container workloads, this presents a wide array of benefits.

On-premises ONTAP: Kubernetes benefits

If you are running an on-premises Kubernetes cluster – whether with a Kubernetes engine (such as Rancher) or one you’ve installed yourself – you probably have to consider providing back-end storage to the container pods. The following sections describe the benefits of using ONTAP on premises next to your on-premises Kubernetes clusters.

Data migration

One of the biggest challenges of large datasets – other than backing them up – is how to move them when you need to leverage new hardware or cloud instances.

Kubernetes applications are not immune to this problem, so when you need to migrate datasets to the cloud or localize the datasets to your compute instances, ONTAP has several options for fast and easy data migration.

- For migrations within the same ONTAP cluster for performance tiering and hardware refreshes, nondisruptive volume moves can greatly reduce the impact to your production Kubernetes instances by reducing the maintenance window and eliminating the need for downtime.
- For migrations to other clusters at remote sites or to the cloud, SnapMirror replication can be used to quickly and efficiently move datasets with minimal downtime. The best part for your Kubernetes workloads is that NetApp Astra Control software provides Kubernetes-native integration for SnapMirror that allows you to migrate the entire Kubernetes stack if necessary.
- If you want to migrate data just to localize reads, consider the ONTAP FlexCache feature*, where you don’t need to migrate any data at all. Simply create a cache attached to an origin dataset and ONTAP automatically fetches the blocks your Kubernetes application requests. Or you can prepopulate selected folders in the dataset to minimize the overall storage footprint needed for your application to finish its work.
- If you need an entire copy of the source data to make changes, you can avoid data migration entirely by standing up a Snapshot backed clone (a FlexClone), which instantly presents an entire copy of a volume that uses space only when your application starts to make changes. As a bonus, if you decide to keep the changes, you can split the clone to a new volume and keep both or delete the original

volume. FlexClone is perfect for backup verification, code version control, and any use case where you want to use a dataset without absorbing massive storage footprint costs.

* Not currently supported with NetApp Astra.

Multiprotocol data access

Flexibility for your Kubernetes storage isn't just being able to serve data anywhere; it's also being able to serve data via any protocol that your application requires. ONTAP provides support for S3, as well as file (NAS) and block (SAN) access to persistent volumes for Kubernetes clusters. Having options for how to present data to containers by using a single storage platform makes life easier for Kubernetes administrators and application owners alike.

Performance

Ask any application owner if they'd like for their application to run faster, and they'll almost always say "yes." ONTAP has many ways to improve overall performance. Some are hardware-based (such as All Flash FAS) and some are software features baked into ONTAP.

Naturally, flash storage is the fastest option for any application – especially if you're able to leverage NVMe-attached SSDs. But for some, flash is cost prohibitive, especially if your datasets aren't consistently in use. There are also hybrid approaches to flash, with a front-end Flash Pool aggregate that can offload random read operations and repetitive random write operations before they make their way to spinning disk.

Of course, SSDs and flash aren't unique to ONTAP; most storage providers have some sort of flash option for their storage systems. However, ONTAP offers features outside of hardware solutions that help to greatly improve overall performance for all applications – including ones running with Kubernetes.

The list is extensive, but here are a handful of things ONTAP provides at a software level, regardless of the type of hardware being used.

- **WAFL® technology.** The Write Anywhere File Layout file system is what makes ONTAP, ONTAP. When new blocks are created, rather than trying to find the perfect place for them, WAFL simply says "put this in the first open spot you see." By avoiding extra overhead in trying to lay data out in a specific pattern, WAFL can often outperform most modern file systems for initial data creation. After that data is created, WAFL has an inode map for reference when those files need to be accessed later for new reads and writes. This architecture has proven to be even better for flash disk than for spinning drives, because it reduces the amount of wear leveling due to the fact that new blocks never overwrite old ones, but are instead written to new blocks by design.
- **NVRAM acceleration.** Incoming writes to ONTAP don't go directly to disk; instead, they are processed in memory and staged in NVRAM while notifying clients that they were acknowledged immediately. After processing, these writes are committed to disk without interaction with the clients. This reduces the amount of latency for workloads by reducing the overall path length for initial write activities.
- **Multithreaded reads and writes.** To avoid CPU bottlenecks, ONTAP leverages multiple CPU cores to process reads and writes. This parallelization helps get work done faster and helps to make sure that all CPUs are being used in the most efficient way possible.
- **NVMe over TCP support*.** If you're using block protocols to host your Kubernetes datasets, you're probably using either FCP or iSCSI LUNs. These have worked fine for the past 20 years or so, but SCSI commands are chatty and inefficient. High-performance applications can get better results by switching over to NVMe transport methods with lower latency and higher throughput from command streamlining, context switch reduction, and increased multithreading. ONTAP supports NVMe over TCP for environments that don't have an existing FC infrastructure in place. If you have existing LUNs, ONTAP also supports converting LUNs to NVMe namespaces without having to migrate any data.

- **Performance enhancements for NAS.** If you're using file protocols (such as NFS) for your Kubernetes workloads, ONTAP has added performance enhancements for that as well. FlexGroup volumes provide massive scale (for both capacity and file counts) across multiple nodes for aggregated throughput and faster file ingest by multithreading write metadata operations. Parallel NFS provides read and write localization to avoid back-end cluster networks. Nconnect support offers additional parallelization by creating multiple TCP sessions for a single NFS mount point to improve queue depth. FlexCache allows you to create multiple sparse caches to help spread out the read workload on a single dataset across more nodes, clusters, or even to the cloud.

* Currently not supported by NetApp Astra Trident.

Scale

Kubernetes clusters can scale dynamically to massive sizes, so you probably want your storage to be able to scale along with them.

ONTAP clusters can scale up to 24 nodes, provide up to 2.4 million TCP connections, 12,000 FlexVol® volumes (nearly 1200PB and 23.9 trillion inodes)*, 20PB single namespace volumes, and millions of IOPS, depending on the hardware and deployment method being used. You can start small with an ONTAP cluster and then nondisruptively scale larger to keep pace with your Kubernetes clusters.

That scale can be extended to archive use cases, where cloud tiering can automatically move blocks to cloud or S3 storage for more cost-effective space utilization. There are many ways to achieve massive scale with your Kubernetes clusters and PVs using ONTAP storage.

* On-premises ONTAP with 100TB FlexVol volumes at 2 billion inodes per volume.

Security

With companies constantly in the news having to report data breaches or ransomware attacks, security has gone from an afterthought in Kubernetes deployments to top of mind. In fact, the US CISA published a Kubernetes hardening guide recently to help administrators better secure their clusters.

Securing Kubernetes is important, but you also need to be able to secure the datasets attached to the Kubernetes clusters. ONTAP provides a slew of features for the security minded. ONTAP encryption passes the standards for NIST FIPS 140-2 and has earned multiple other security certifications and compliance qualifications.

- **At-rest encryption.** You can encrypt your data at the hardware layer by using NetApp Storage Encryption (NSE) or self-encrypting drives (SED), or at the ONTAP software layer by using aggregate-level or volume-level AES-256 bit encryption via an onboard key manager in ONTAP or off-box key managers.
- **In-flight encryption.** ONTAP supports several ways to encrypt data in flight with the same stringent focus on security as the at-rest options. For NAS workloads, ONTAP supports Kerberos encryption for SMB and NFS, IPSec for any workload traversing an ethernet connection, TLS 1.2 encryption for FlexCache and SnapMirror technologies, and TLS 1.3 support for cluster logins.
- **Autonomous ransomware protection*.** Although ONTAP does not natively stop ransomware attacks (most of those attacks are performed from authenticated users being social engineered or phished), it can detect data anomalies to alert administrators and take an automatic Snapshot copy. And automatic ransomware protection allows quick recovery time without needing to pay a ransom.
- **Multifactor authentication**. An ONTAP admin account is a key attack vector, because a bad actor can access a cluster and start deleting volumes, Snapshot copies, and much more. Multifactor authentication provides an additional layer of protection from attackers getting admin access to a cluster hosting Kubernetes datasets.
- **Multi-admin verification.** If an attacker gets access to an admin account in ONTAP, usually all bets are off. However, with the new multi-admin verification feature in ONTAP, you can specify lists of

tasks that require multiple sign-offs to perform, such as volume or snapshot deletions. This is an especially critical feature in combatting ransomware attacks.

- **Cloud Insights and Cloud Secure.** Cloud Insights and Cloud Secure integrate tightly with ONTAP to offer deep analytics of data anomalies and visibility into the source of attacks, and they leverage machine learning to deliver fewer false positives when alerting to a ransomware event. Plus, Cloud Insights provides the ability to plug into your Kubernetes clusters for log retrieval and monitoring.

* Currently not supported with NetApp Astra Trident.

Storage efficiencies

Applications that require large datasets for operation often have a hidden secret – much of the data in those datasets is either duplicated or isn't used regularly. In those cases, you may be paying extra to keep that data around when it's not even being used. Storage real estate is pricey, especially when you're dealing with SSDs and on-premises power consumption.

ONTAP storage efficiencies offer a way to reduce those costs by deduplicating and compressing data, both in real time and after placement. They also offer automatic cloud tiering with FabricPool to move underutilized data to more cost-effective cloud or object storage systems, while maintaining the ability to automatically rehydrate those datasets when the application requests them again. In addition, space efficient Snapshot copies enable building a finer granular low-RPO backup and DR solution for Kubernetes apps. Plus, in the case of AFF systems, NetApp provides guaranteed efficiencies for specific dataset types.

ONTAP for Kubernetes benefits: In the cloud

If you're looking to move your Kubernetes workloads to the cloud, or are already running them there, ONTAP has a cloud-native presence that brings a rich and growing feature set that enhances your Kubernetes experience with secure, performant storage that offers robust data protection.

ONTAP in the cloud: Options

If you want to run ONTAP in the cloud, you have two distinct options across all of the major cloud providers.

Self-managed ONTAP with Cloud Volumes ONTAP and Amazon FSx

You can run a virtualized instance of ONTAP in the cloud where you have control over the software aspects of ONTAP (including access to the CLI) while not having to deal with the hardware underpinnings such as disks, nodes, and aggregates. For Azure and Google Cloud, Cloud Volumes ONTAP would be your choice, while Amazon FSx would be the option in AWS.

Use a self-managed ONTAP offering if you require:

- More control over the ONTAP configuration
- Feature sets closer to what on-premises ONTAP offers
- Access to CLI

ONTAP as a service in the cloud

If you don't need control over the storage system and simply want to provision NAS storage to your clients or have them self-provision their storage, then you should consider NetApp Cloud Volumes Service in Google Cloud and AWS, or Azure NetApp Files for a more hands-off approach to infrastructure management. ONTAP as a service in the cloud presents a limited feature set in ONTAP with an easy-to-use GUI management system (or access to REST API) to streamline basic volume presentation to end users.

ONTAP in the cloud: Kubernetes benefits

Running ONTAP in the cloud means that you get most of the benefits of on-premises ONTAP without the extra overhead of managing hardware or even configuring those instances (depending on which cloud offering you've chosen).

With Cloud Volumes ONTAP and Amazon FSx, you get virtually all of the same benefits as listed in the section titled "ONTAP for Kubernetes benefits: On premises."

With Cloud Volumes Service and Azure NetApp Files, the feature set and configuration options are limited to what is currently supported by the Cloud Volumes Service provider, so you need to check the documentation for your selected cloud provider, but the general concepts are covered in the following subsections.

Data migration: Cloud Volumes Service

Cloud Volumes Service currently only provides file-based data migration, which can be in the form of NetApp XCP or NetApp Cloud Sync. Although Cloud Volumes Service supports NetApp SnapMirror to replicate between instances, it does not currently support replication from non-Cloud Volumes Services ONTAP instances, such as Cloud Volumes ONTAP or on-premises ONTAP. However, NetApp's file-based migration tools are fast and efficient and work great with any dataset, including those backed by Kubernetes.

For more information on cloud data migration, download the free NetApp guide, Plan Your Lift & Shift Cloud Migration Strategy.

Data protection

NetApp ONTAP offers enterprise-class data protection features such as instant Snapshot copies and fast, efficient replication via NetApp SnapMirror. Cloud Volumes Service offers both of these features, and when you combine those with the new SnapMirror support in Astra Control, you get a complete data protection strategy for your Kubernetes cluster, from data volumes to Kubernetes configuration files, secrets, and more.

For more information on Kubernetes data protection using SnapMirror:

- What customers say about ONTAP SnapMirror coming to Astra Control
- Astra adds AWS support and SnapMirror based DR for Kubernetes workloads
- Tech ONTAP Podcast Episode 339 - NetApp Astra Control Updates - August 2022

Multiprotocol data access: Cloud Volumes Service

Cloud Volumes Service currently only provides file-based data access (NFS, SMB, and a dual-protocol mixture of NFS and SMB). As a result, if your Kubernetes workloads use NAS protocols to present persistent volume claims, then Cloud Volumes Service is a great fit, especially if your Kubernetes cluster runs in a cloud-native service such as GKE, AKS, or EKS.

Performance: Cloud Volumes Service

Cloud Volumes Service is designed to remove the administrative day-to-day tasks for storage provisioning, including how to optimize performance. If you want to get the best possible results for your workloads. simply select the Performance instance tier and Cloud Volumes Service does the rest on the back end. If you don't require high performance for your Kubernetes workloads, then select the less expensive Standard option and start provisioning your PVs.

Scale: Cloud Volumes Service

With on-premises ONTAP, you can scale clusters by adding extra nodes, using bigger platforms, adding data LIFs, and so on. Cloud Volumes Service doesn't incorporate hardware management. Instead, scale is achieved by adding new volumes to your Kubernetes workloads. In some deployments, you can even leverage FlexGroup volumes for additional scale of capacity and performance. See the documentation for your Cloud Volumes Service deployment for more information about what is and is not supported.

Security: Cloud Volumes Service

Because you don't manage Cloud Volumes Service instances, security is already baked into the service. You can configure additional security features, such as Kerberos for in-flight NFS encryption for your Kubernetes datasets, but ultimately, you'll get a deployment that can be trusted to be secure on the back end with a combination of the cloud provider's security options and the preconfigured ONTAP security features. For an example of the security offered in Cloud Volumes Service, read the technical report, TR-4918: Security overview - NetApp Cloud Volumes Service in Google Cloud.

NetApp Astra: Kubernetes benefits

NetApp Astra is has an application suite targeted at Kubernetes workloads that can make management of external storage systems simpler and help encourage a DevOps workflow by moving the responsibilities for application data management away from storage administrators to the application owners. NetApp Astra can also back up entire Kubernetes clusters, from stateful datasets to the Kubernetes configuration files and secrets.

The NetApp Astra suite includes the following.

NetApp Astra Trident, a CSI driver for easily provisioning external storage for Kubernetes clusters and containers and enabling simpler DevOps workflows by allowing end users to provision their own storage.

NetApp Astra Control Center, a self-managed, application-aware data management offering designed for use with Kubernetes clusters. Astra Control Center allows Kubernetes administrators to easily back up and restore Kubernetes clusters and stateful datasets, migrate workloads, and create instant, space-efficient clones for dev/test workflows.

NetApp Astra Control Service, a NetApp managed, cloud-resident "as a service" offering of NetApp Astra Control. Astra Control Service offers data protection and migration services for your cloud-based Kubernetes deployments, whether they live in Azure, AWS, or GC.

For more detailed information about NetApp Astra, refer to the following links:

- Astra Control on Azure Kubernetes Service - YouTube
- Simplified Kubernetes data management with Astra Control
- NetApp beefs up its Astra Control plane for Kubernetes apps
- NetApp Astra and Kubernetes – Tech ONTAP Podcast playlist

Where to find additional information

To learn more about the information that is described in this document, check out the following documents and/or websites:

- TR-4918: Security overview: NetApp Cloud Volumes Service in Google Cloud
  docs.netapp.com/us-en/netapp-solutions/ehc/ncvs/ncvs-gc-overview.html
- NetApp Astra Control: A deep dive into the use cases
  bluexp.netapp.com/blog/blg-netapp-astra-control-a-deep-dive-into-the-use-cases

- What customers say about ONTAP SnapMirror coming to Astra Control
  bluexp.netapp.com/blog/astra-blg-what-customers-are-saying-about-ontap-snapmirror-coming-to-astra-control
- Business continuity for your Kubernetes applications with NetApp Astra Control
  bluexp.netapp.com/blog/astra-blg-business-continuity-for-your-kubernetes-applications-with-netapp-astra-control
- Kubernetes Storage: An In-Depth Look
  bluexp.netapp.com/blog/cvo-blg-kubernetes-storage-an-in-depth-look
- NetApp Trident GitHub
  github.com/NetApp/trident
- NetApp Tech ONTAP Podcast
  techontappodcast.com/
- NetApp product documentation
  https://www.netapp.com/support-and-training/documentation/

Version history

Version	Date	Document version history
Version 1.0	January 2023	Initial release.

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright information

Trademark information

WP-7360--0123

NetApp