



Concepts

Astra Control Center

NetApp
August 11, 2025

Table of Contents

Concepts	1
Architecture and components	1
Capabilities	1
Architecture	1
Deployment models	2
For more information	3
Data protection	3
Snapshots, backups, and protection policies	3
Clones	4
Replication between storage backends	4
Backups, snapshots, and clones with an expired license	6
Licensing	6
Evaluation licenses and full licenses	7
License expiration	7
How license consumption is calculated	7
Find more information	7
App management	7
Storage classes and persistent volume size	9
Overview	9
Storage classes	9
User roles and namespaces	10
User roles	10
Namespaces	10
Find more information	10

Concepts

Architecture and components

Astra Control is a Kubernetes application data lifecycle management solution that simplifies operations for stateful applications, and helps you store, protect, and move your Kubernetes workloads across hybrid environments.

Capabilities

Astra Control offers critical capabilities for Kubernetes application data lifecycle management:

Store:

- Dynamic storage provisioning for containerized workloads
- In-flight encryption of data from container to persistent volumes
- Cross-region, cross-zone replication

Protect:

- Automated discovery and application-aware protection of an entire application and its data
- Instant recovery of an application from any snapshot version based on your organization's needs
- Fast failover across zones, regions and cloud providers

Move:

- Complete application and data mobility within and between Kubernetes clusters and clouds
- Instant clones of entire applications and data
- One click migration of applications through consistent web UI and API

Architecture

The architecture of Astra Control enables it to provide advanced data management capabilities that enhance both the functionality and availability of Kubernetes applications, simplifies the management, protection, and movement of containerized workloads across public clouds and on-prem environments, and provides automation capabilities through its REST API and SDK, enabling programmatic access for seamless integration with existing workflows.

Astra Control is Kubernetes-native, enabling data protection workflows that utilize custom resources while staying backward-compatible with the existing API and SDK. Kubernetes-native data protection offers significant advantages; by seamlessly integrating with Kubernetes APIs and resources, data protection can become an inherent part of the application lifecycle through an organization's existing CI/CD and/or GitOps tools.

Astra Control is built on four complementary components:

- **Astra Control:** Astra Control is the centralized management service for all managed clusters, providing orchestrated workloads for application protection and mobility for on-premises as well as the following capabilities:

- Combined view of multiple clusters
 - Protection of orchestrated workflows
 - Granular resource visualization and selection
- **Astra Connector:** Astra Connector teams with Astra Control to provide a secure connection to each managed cluster, offering local execution of scheduled operations regardless of connection status as well as the following capabilities:
 - Local execution of scheduled operations regardless of connection status
 - Local operations that distribute and optimize system resource usage of Astra across clusters
 - Local installation that enables least privilege access to the cluster for improved security
- **Astra Control Provisioner:** Astra Control Provisioner delivers core CSI provisioning functionality and advanced storage management capabilities for added security and disaster recovery configuration, as well as the following capabilities:
 - Dynamic storage provisioning for containerized workloads
 - Advanced storage management:
 - In-flight encryption of data from container to PV
 - SnapMirror Cloud functionality with cross-region, cross-zone replication
- **Astra Custom resources:** Custom resources used on each cluster provide a Kubernetes-native approach to running operations locally, simplifying integration with other Kubernetes-friendly tooling and automation as well as providing the following capabilities:
 - Direct ecosystem tool integration and automation workflows
 - Lower-level primitives that enable custom workflows

Deployment models

Astra Control is available in a single deployment model.

Astra Control Center: Self-managed software that provides application-aware data management of Kubernetes clusters running in your on-premises environment. Astra Control Center can also be installed on multiple cloud provider environments with a NetApp Cloud Volumes ONTAP storage backend.

[Astra Control Center documentation](#)

Astra Control Center	
How is it offered?	As software that you can download, install, and manage
Where is it hosted?	On your own Kubernetes cluster
How is it updated?	You manage any updates
What are the supported Kubernetes distributions?	<ul style="list-style-type: none"> • Azure Kubernetes Service on Azure Stack HCI • Google Anthos • Kubernetes (Upstream) • Rancher Kubernetes Engine (RKE) • Red Hat OpenShift Container Platform

	Astra Control Center
What are the supported storage backends?	<ul style="list-style-type: none"> NetApp ONTAP AFF and FAS systems NetApp ONTAP Select Cloud Volumes ONTAP Longhorn

For more information

- [Astra Control Center documentation](#)
- [Astra Trident documentation](#)
- [Astra Control API](#)
- [Cloud Insights documentation](#)
- [ONTAP documentation](#)

Data protection

Learn about the available types of data protection in Astra Control Center, and how best to use them to protect your apps.

Snapshots, backups, and protection policies

Both snapshots and backups protect the following types of data:

- The application itself
- Any persistent data volumes associated with the application
- Any resource artifacts belonging to the application

A *snapshot* is a point-in-time copy of an app that's stored on the same provisioned volume as the app. They are usually fast. You can use local snapshots to restore the application to an earlier point in time. Snapshots are useful for fast clones; snapshots include all of the Kubernetes objects for the app, including configuration files. Snapshots are useful for cloning or restoring an app within the same cluster.

A *backup* is based on a snapshot. It is stored in the external object store, and because of this, can be slower to take compared to local snapshots. You can restore an app backup to the same cluster, or you can migrate an app by restoring its backup to a different cluster. You can also choose a longer retention period for backups. Because they are stored in the external object store, backups generally offer you better protection than snapshots in cases of server failure or data loss.

A *protection policy* is a way to protect an app by automatically creating snapshots, backups, or both according to a schedule that you define for that app. A protection policy also enables you to choose how many snapshots and backups to retain in the schedule, and set different schedule granularity levels. Automating your backups and snapshots with a protection policy is the best way to ensure each app is protected according to the needs of your organization and service level agreement (SLA) requirements.

 *You can't be fully protected until you have a recent backup.* This is important because backups are stored in an object store away from the persistent volumes. If a failure or accident wipes out the cluster and its associated persistent storage, then you need a backup to recover. A snapshot would not enable you to recover.

Immutable backups

An immutable backup is a backup that cannot be changed or deleted during a specified period. When you create an immutable backup, Astra Control checks to ensure that the bucket you are using is a write once read many (WORM) bucket, and if so, ensures that the backup is immutable from within Astra Control. Astra Control Center supports creating immutable backups with the following platforms and bucket types:

- Amazon Web Services using an Amazon S3 bucket with S3 Object Lock configured
- NetApp StorageGRID using an S3 bucket with S3 Object Lock configured

Note the following when working with immutable backups:

- If you back up to a WORM bucket in an unsupported platform or to an unsupported bucket type, you might get unpredictable results, such as backup deletion failing even if the retention time has elapsed.
- Astra Control does not support data lifecycle management policies or manual deletion of objects on the buckets you use with immutable backups. Make sure that your storage backend is not configured to manage the lifecycle of Astra Control snapshots or backed up data.

Clones

A *clone* is an exact duplicate of an app, its configuration, and its persistent data volumes. You can manually create a clone on either the same Kubernetes cluster or on another cluster. Cloning an app can be useful if you need to move applications and storage from one Kubernetes cluster to another.

Replication between storage backends

Using Astra Control, you can build business continuity for your applications with a low-RPO (Recovery Point Objective) and low-RTO (Recovery Time Objective) using asynchronous replication capabilities of NetApp SnapMirror technology. Once configured, this enables your applications to replicate data and application changes from one storage backend to another, on the same cluster or between different clusters.

You can replicate between two ONTAP SVMs on the same ONTAP cluster or on different ONTAP clusters.

Astra Control asynchronously replicates app snapshot copies to a destination cluster. The replication process includes data in the persistent volumes replicated by SnapMirror and the app metadata protected by Astra Control.

App replication is different from app backup and restore in the following ways:


- **App replication:** Astra Control requires the source and destination Kubernetes clusters (which can be the same cluster) to be available and managed with their respective ONTAP storage backends configured to enable NetApp SnapMirror. Astra Control takes the policy-driven application snapshot and replicates it to the destination storage backend. NetApp SnapMirror technology is used to replicate the persistent volume data. To fail over, Astra Control can bring the replicated app online by recreating the app objects on the destination Kubernetes cluster with the replicated volumes on the destination ONTAP cluster. Because the persistent volume data is already present on the destination ONTAP cluster, Astra Control can offer quick recovery times for failover.

- **App backup and restore:** When backing up applications, Astra Control creates a snapshot of the app data and stores it in an object storage bucket. When a restore is needed, the data in the bucket must be copied to a persistent volume on the ONTAP cluster. The backup/restore operation does not require the secondary Kubernetes/ONTAP cluster to be available and managed, but the additional data copy can result in longer restore times.


To learn how to replicate apps, refer to [Replicate apps to a remote system using SnapMirror technology](#).

The following images show the scheduled backup and restore process compared to the replication process.

The backup process copies data to S3 buckets and restores from S3 buckets:



On the other hand, replication is done by replicating to ONTAP and then a failover creates the Kubernetes resources:



Backups, snapshots, and clones with an expired license

If your license expires, you can add a new application or perform application protection operations (such as snapshots, backups, clones, and restore operations) only if the application you are adding or protecting is another Astra Control Center instance.

Licensing

When you deploy Astra Control Center, it is installed with an embedded 90-day evaluation license for 4,800 CPU units. If you need more capacity or a longer evaluation period, or want to upgrade to a full license, you can obtain a different evaluation license or full license from NetApp.

You obtain a license in one of the following ways:

- If you are evaluating Astra Control Center and need different evaluation terms than what is included in the embedded evaluation license, contact NetApp to request a different evaluation license file.
- [If you already purchased Astra Control Center, generate your NetApp license file \(NLF\)](#) by logging in to the NetApp Support Site and navigating to your software licenses under the Systems menu.

For details about licenses needed for ONTAP storage backends, refer to [supported storage backends](#).



Make sure that your license enables at least as many CPU units as you need. If the number of CPU units that Astra Control Center is currently managing exceeds the available CPU units in the new license being applied, you'll not be able to apply the new license.

Evaluation licenses and full licenses

An embedded evaluation license is provided with a new Astra Control Center installation. An evaluation license enables the same capabilities and features as a full license for a limited (90 day) period. After the evaluation period, a full license is required to continue with full functionality.

License expiration

If the active Astra Control Center license expires, UI and API functionality for the following features are unavailable:

- Manual local snapshots and backups
- Scheduled local snapshots and backups
- Restoring from a snapshot or backup
- Cloning from a snapshot or current state
- Managing new applications
- Configuring replication policies

How license consumption is calculated

When you add a new cluster to Astra Control Center, it doesn't count toward consumed licenses until at least one application running on the cluster is managed by Astra Control Center.

When you start managing an app on a cluster, all of that cluster's CPU units are included in the Astra Control Center license consumption, except Red Hat OpenShift cluster node CPU units that are reported by a using the label `node-role.kubernetes.io/infra: ""`.



Red Hat OpenShift infrastructure nodes do not consume licenses in Astra Control Center. To mark a node as an infrastructure node, apply the label `node-role.kubernetes.io/infra: ""` to the node.


Find more information

- [Add a license when you first set up Astra Control Center](#)
- [Update an existing license](#)


App management

When Astra Control discovers your clusters, the apps on those clusters are unmanaged until you choose how you want to manage them. A managed application in Astra Control can be any of the following:


- A namespace, including all resources in that namespace



- An individual application deployed within one or more namespaces (helm3 is used in this example)



- A group of resources that are identified by a Kubernetes label within one or more namespaces



Storage classes and persistent volume size

Astra Control Center supports NetApp ONTAP and Longhorn as storage backends.

Overview

Astra Control Center supports the following:

- **Storage classes backed by ONTAP storage:** If you are using an ONTAP backend, Astra Control Center offers the ability to import the ONTAP backend to report monitoring information.
- **CSI-based storage classes backed by Longhorn:** You can use Longhorn with the Longhorn Container Storage Interface (CSI) driver.



Storage classes should be [configured](#) using Astra Control Provisioner.

Storage classes

When you add a cluster to Astra Control Center, you're prompted to select one previously configured storage class on that cluster as the default storage class. This storage class will be used when no storage class is specified in a persistent volume claim (PVC). The default storage class can be changed at any time within Astra Control Center and any storage class can be used at any time by specifying the name of the storage class within the PVC or Helm chart. Ensure that you have only a single default storage class defined for your Kubernetes cluster.

User roles and namespaces

Learn about user roles and namespaces in Astra Control, and how you can use them to control access to resources in your organization.

User roles

You can use roles to control the access users have to resources or capabilities of Astra Control. The following are the user roles in Astra Control:

- A **Viewer** can view resources.
- A **Member** has Viewer role permissions and can manage apps and clusters, unmanage apps, and delete snapshots and backups.
- An **Admin** has Member role permissions and can add and remove any other users except the Owner.
- An **Owner** has Admin role permissions and can add and remove any user accounts.

You can add constraints to a Member or Viewer user to restrict the user to one or more [Namespaces](#).

Namespaces

A namespace is a scope you can assign to specific resources within a cluster that is managed by Astra Control. Astra Control discovers a cluster's namespaces when you add the cluster to Astra Control. Once discovered, the namespaces are available to assign as constraints to users. Only members that have access to that namespace are able to use that resource. You can use namespaces to control access to resources using a paradigm that makes sense for your organization; for example, by physical regions or divisions within a company. When you add constraints to a user, you can configure that user to have access to all namespaces or only a specific set of namespaces. You can also assign namespace constraints using namespace labels.

Find more information

[Manage local users and roles](#)

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.