



Concepts

Astra Control Service

NetApp
March 11, 2024

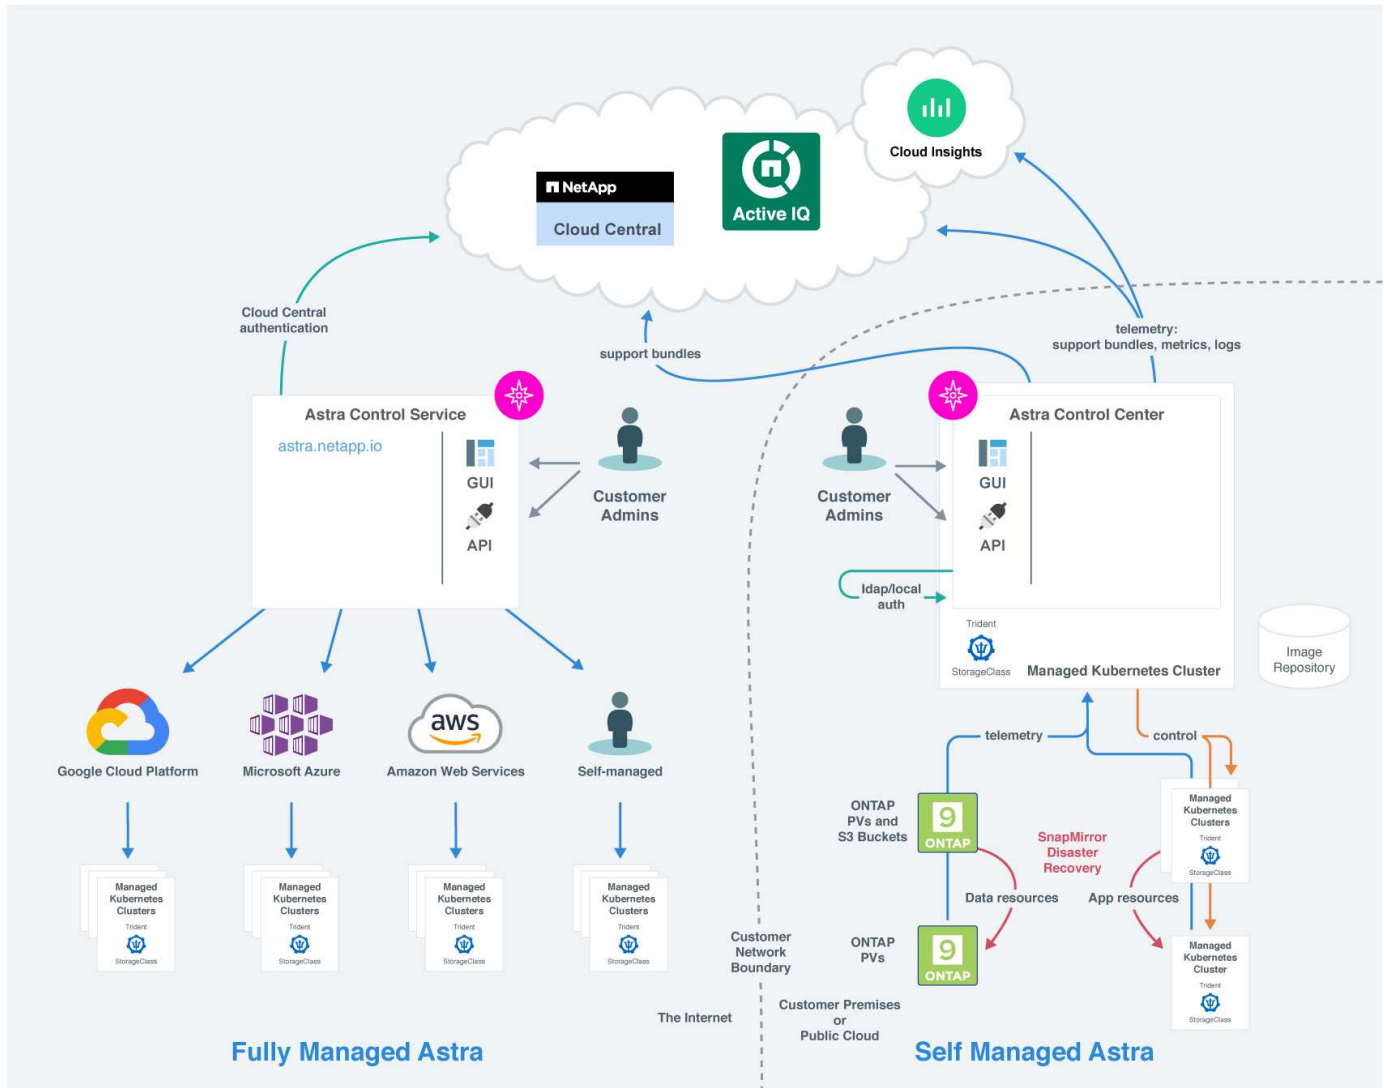
Table of Contents

- Concepts 1
 - Architecture and components 1
 - Data protection 2
 - Storage classes and performance for AWS clusters 3
 - Storage classes and PV size for AKS clusters 4
 - Service type, storage classes, and PV size for GKE clusters 5
 - App management 8
 - User roles and namespaces 10

Concepts

Architecture and components

Here is an overview of the various components of the Astra Control environment.



Astra Control components

- **Kubernetes clusters:** Kubernetes is a portable, extensible, open-source platform for managing containerized workloads and services, that facilitates both declarative configuration and automation. Astra provides management services for applications hosted in a Kubernetes cluster.
- **Astra Trident:** As a fully supported open source storage provisioner and orchestrator maintained by NetApp, Astra Trident enables you to create storage volumes for containerized applications managed by Docker and Kubernetes. When deployed with Astra Control Center, Astra Trident includes a configured ONTAP storage backend.
- **Cloud Insights:** A NetApp cloud infrastructure monitoring tool, Cloud Insights enables you to monitor performance and utilization for your Kubernetes clusters managed by Astra Control Center. Cloud Insights correlates storage usage to workloads. When you enable the Cloud Insights connection in Astra Control

Center, telemetry information shows in Astra Control Center UI pages.

Astra Control interfaces

You can complete tasks using different interfaces:

- **Web user interface (UI):** Both Astra Control Service and Astra Control Center use the same web-based UI where you can manage, migrate and protect apps. Use the UI also to manage user accounts and configuration settings.
- **API:** Both Astra Control Service and Astra Control Center use the same Astra Control API. Using the API, you can perform the same tasks that you would using the UI.

Astra Control Center also enables you to manage, migrate, and protect Kubernetes clusters running within VM environments.

For more information

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

Data protection

Learn about the available types of data protection in Astra Control Service, 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.



If you perform a snapshot or backup, but the operation fails with the error "The resource wasn't created because of an internal server issue", check to make sure the storage backend you are using has the correct drivers installed. Some storage backends need Container Storage Interface (CSI) drivers, while others need an external snapshot controller.

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 Service supports creating immutable backups with the following platforms and bucket types:

- Amazon Web Services using an Amazon S3 bucket with S3 Object Lock configured
- Microsoft Azure using an Azure bucket with a retention policy configured
- Google Kubernetes Engine (GKE) using a Google Cloud Storage bucket with a retention policy 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.

Storage classes and performance for AWS clusters

Astra Control Service can use Amazon Elastic Block Store (EBS), Amazon FSx for NetApp ONTAP, or NetApp Cloud Volumes ONTAP as the storage backend for Amazon Elastic Kubernetes Service (EKS) clusters.

Amazon Elastic Block Store (EBS)

Your clusters can use Container Storage Interface (CSI) drivers to interface with EBS. When you use EBS as the storage backend for EKS clusters, you can configure some storage class parameters. For more information about what the parameters mean and how to configure them, refer to [the Kubernetes documentation](#).

You can use several different types of volumes with EBS:

- Solid state drives (SSD)
- Hard disk drives (HDD)
- Previous generation

For more information on each type of volume and their performance, refer to [the Amazon EBS documentation](#). For pricing information, refer to [Amazon EBS pricing](#).

Amazon FSx for NetApp ONTAP

When you use FSx for NetApp ONTAP as the storage backend for AWS clusters, I/O performance depends on the configuration of the filesystem and the characteristics of your workloads. For specific information on FSx for NetApp ONTAP performance, refer to [Amazon FSx for NetApp ONTAP performance](#). For pricing information, refer to [Amazon FSx for NetApp ONTAP Pricing](#).

NetApp Cloud Volumes ONTAP

For specific information on configuring NetApp Cloud Volumes ONTAP, including performance recommendations, visit the [NetApp Cloud Volumes ONTAP documentation](#).

Storage classes and PV size for AKS clusters

Astra Control Service supports Azure NetApp Files, Azure managed disks, or NetApp Cloud Volumes ONTAP as the storage backend for Azure Kubernetes Service (AKS) clusters.

Azure NetApp Files

Astra Control Service supports Azure NetApp Files as the storage backend for Azure Kubernetes Service (AKS) clusters. You should understand how choosing a storage class and persistent volume size can help you meet your performance objectives.

Service levels and storage classes

Azure NetApp Files supports three service levels: Ultra storage, Premium storage, and Standard storage. Each of these service levels are designed for different performance needs:

Ultra storage

Provides up to 128 MiB/s of throughput per 1 TiB.

Premium storage

Provides up to 64 MiB/s of throughput per 1 TiB.

Standard storage

Provides up to 16 Mib/s of throughput per 1 TiB.

These service levels are an attribute of a capacity pool. You need to set up a capacity pool for each service level that you want to use with your Kubernetes clusters. [Learn how to set up capacity pools.](#)

Astra Control Service uses these service levels as storage classes for your persistent volumes. When you add Kubernetes clusters to Astra Control Service, you're prompted to choose either Ultra, Premium, or Standard as the default storage class. The names of the storage classes are *netapp-anf-perf-ultra*, *netapp-anf-perf-premium*, and *netapp-anf-perf-standard*.

[Learn more about these service levels in the Azure NetApp Files docs.](#)

Persistent volume size and performance

As described above, the throughput for each service level is per 1 TiB of provisioned capacity. That means larger volumes provide better performance. So you should take both capacity and performance needs into consideration when provisioning volumes.

Minimum volume size

Astra Control Service provisions persistent volumes using a minimum volume size of 100 GiB, even if the PVC asks for a smaller volume size. For example, if the PVC in a Helm chart asks for 6 GiB, Astra Control Service automatically provisions a 100 GiB volume.

Application backups

If you back up an application that resides on Azure NetApp Files storage, Astra Control Service automatically temporarily expands the capacity pool. After the backup is complete, Astra Control Service shrinks the capacity pool to its previous size. Depending on your Azure subscription, you might incur storage charges when this happens. You can see a history of capacity pool resize events in the **Activity** page event log.

If the capacity pool exceeds the maximum size allowed by the Azure subscription during the resize operation, the backup operation fails, and a warning is triggered from the Azure API.

Azure managed disks

Astra Control Service can use Container Storage Interface (CSI) drivers to interface with Azure Managed Disks as a storage backend. This service provides block-level storage that is managed by Azure.

[Learn more about Azure managed disks.](#)

NetApp Cloud Volumes ONTAP

For specific information on configuring NetApp Cloud Volumes ONTAP, including performance recommendations, visit the [NetApp Cloud Volumes ONTAP documentation](#).

Service type, storage classes, and PV size for GKE clusters

Astra Control Service supports NetApp Cloud Volumes Service for Google Cloud, Google Persistent Disk, or NetApp Cloud Volumes ONTAP as the storage backend options for persistent volumes.

Cloud Volumes Service for Google Cloud

Astra Control Service can use Cloud Volumes Service for Google Cloud as the storage backend for persistent volumes. You should understand how choosing a service type, storage class, and persistent volume size can help you meet your performance objectives.

Overview

Cloud Volumes Service for Google Cloud provides two service types: *CVS* and *CVS-Performance*. These service types are supported in specific Google Cloud regions. [Go to NetApp BlueXP Global Regions Maps](#) to identify the service type that's supported in the Google Cloud region where your clusters reside.

If your Kubernetes clusters must reside in a specific region, then you'll be using the service type supported in that region.

But if you have the flexibility to choose between Google Cloud regions, then we recommend the following based on your performance requirements:

- For K8s applications that have medium-to-high performance storage needs, choose a Google Cloud region that supports CVS-Performance and use the Premium or Extreme storage class. Such workloads include AI/ML pipelines, CI/CD pipelines, media processing, and databases including relational, noSQL, time series, etc.
- For K8s applications that have low-to-medium storage performance needs (web apps, general purpose file storage, etc.), choose a Google Cloud region that supports either CVS or CVS-Performance, with the Standard storage class.



If you use the CVS service type with Astra Trident 23.01 or later, you need to configure storage pools before you can provision volumes. If you provision volumes with no storage pools configured, volume provisioning will fail. Refer to the [Cloud Volumes Service documentation](#) for more information about creating volumes.

The following table provides a quick comparison of the information described on this page.

Service type	Use case	Supported regions	Storage classes	Min volume size
CVS-Performance	Apps with medium-to-high storage performance needs	View supported Google Cloud regions	<ul style="list-style-type: none">• netapp-cvs-perf-standard• netapp-cvs-perf-premium• netapp-cvs-perf-extreme	100 GiB
CVS	Apps with low-to-medium storage performance needs	View supported Google Cloud regions	netapp-cvs-standard	300 GiB

CVS-Performance service type

Learn more about the CVS-Performance service type before you choose a storage class and create persistent volumes.

Storage classes

Three service levels are supported with the CVS-Performance service type: Standard, Premium, and Extreme. When you add a cluster to Astra Control Service, you're prompted to choose either Standard, Premium, or Extreme as the default storage class for persistent volumes. Each of these service levels are designed for different capacity and bandwidth needs.

The names of the storage classes are *netapp-cvs-perf-standard*, *netapp-cvs-perf-premium*, and *netapp-cvs-perf-extreme*.

[Learn more about these service levels in the Cloud Volumes Service for Google Cloud docs.](#)

Persistent volume size and performance

[As the Google Cloud docs explain](#), the allowed bandwidth for each service level is per GiB of provisioned capacity. That means larger volumes will provide better performance.

Be sure to read through the Google Cloud page linked to above. It includes cost comparisons and examples that can help you better understand how to couple a service level with volume size to meet your performance objectives.

Minimum volume size

Astra Control Service provisions persistent volumes using a minimum volume size of 100 GiB with the CVS-Performance service type, even if the PVC requests a smaller volume size. For example, if the PVC in a Helm chart asks for 6 GiB, Astra Control Service automatically provisions a 100 GiB volume.

CVS service type

Learn more about the CVS service type before you choose a storage class and create persistent volumes.

Storage class

One service level is supported with the CVS service type: Standard. When you manage clusters in regions where the CVS service type is supported, Astra Control Service uses the Standard service level as the default storage class for persistent volumes. The storage class is named *netapp-cvs-standard*.

[Learn more about the Standard service level in the Cloud Volumes Service for Google Cloud docs.](#)

Persistent volume size and performance

The allowed bandwidth for the CVS service type is per GiB of provisioned capacity. That means larger volumes will provide better performance.

Minimum volume size

Astra Control Service provisions persistent volumes using a minimum volume size of 300 GiB with the CVS service type, even if the PVC asks for a smaller volume size. For example, if 20 GiB is requested, Astra Control Service automatically provisions a 300 GiB volume.

Due to a limitation, if a PVC requests a volume between 700-999 GiB, Astra Control Service automatically provisions a volume size of 1000 GiB.

Google Persistent Disk

Astra Control Service can use Container Storage Interface (CSI) drivers to interface with Google Persistent Disk as a storage backend. This service provides block-level storage that is managed by Google.

[Learn more about Google Persistent Disk.](#)

[Learn more about different performance levels of Google Persistent Disks.](#)


NetApp Cloud Volumes ONTAP

For specific information on configuring NetApp Cloud Volumes ONTAP, including performance recommendations, visit the [NetApp Cloud Volumes ONTAP documentation](#).

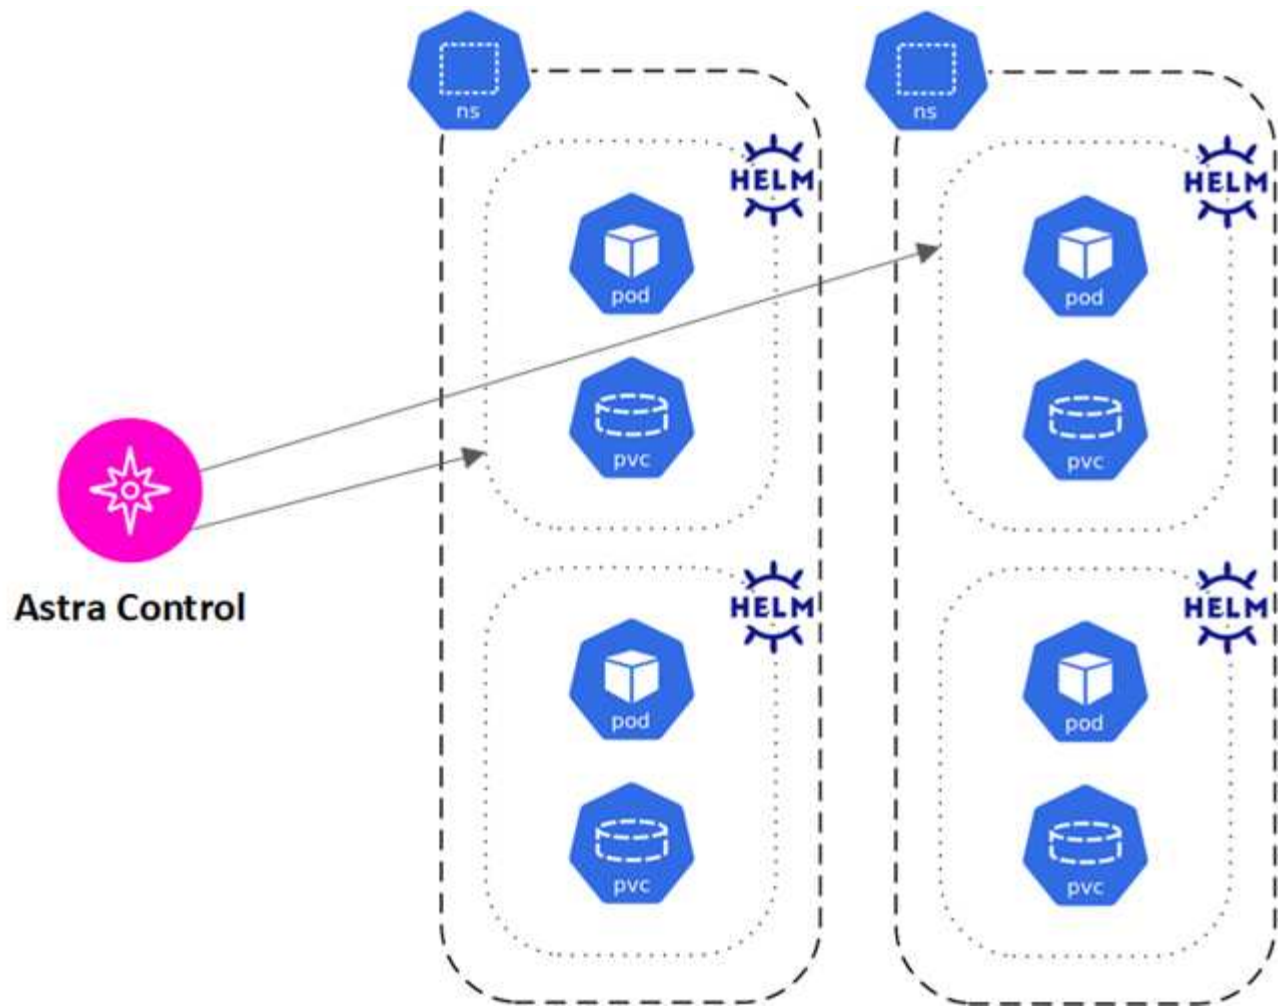
App management

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:

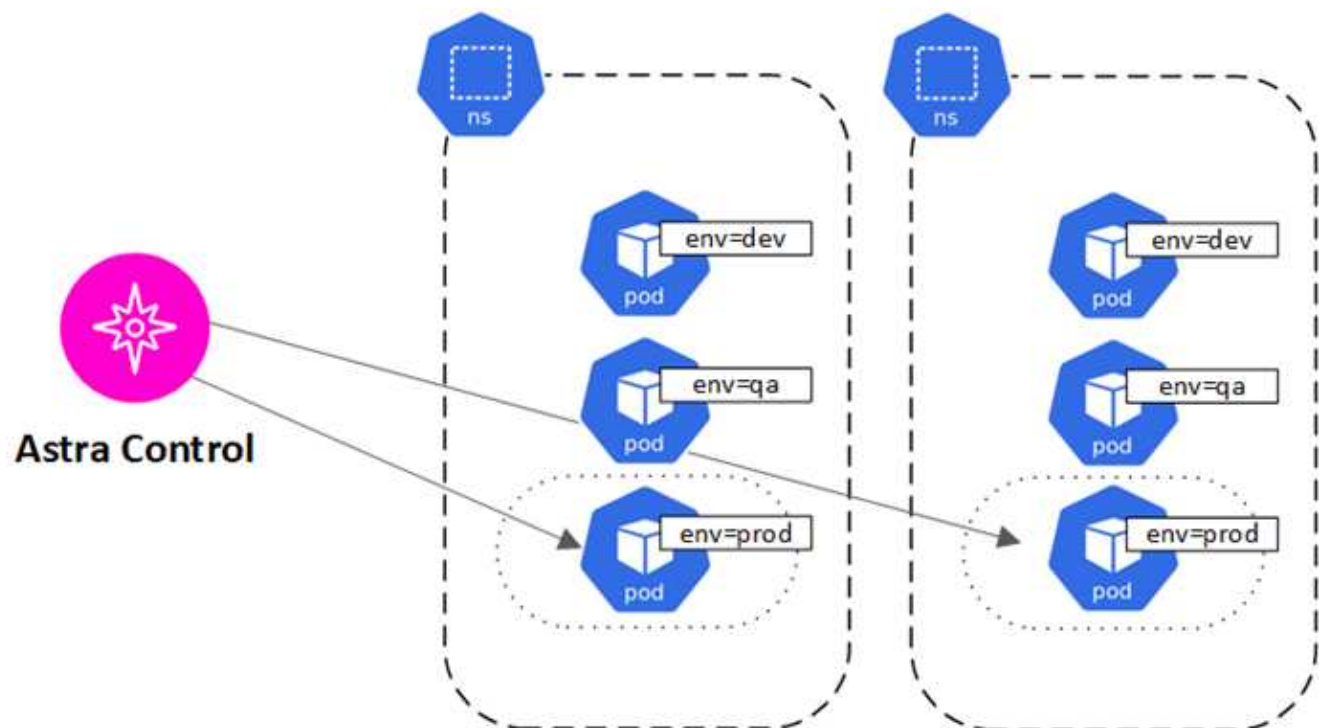
- A namespace, including all resources in that namespace



- An individual application deployed within one or more namespaces (Helm 3 is used in this example)



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



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:

- An **Owner** has Admin permissions and can delete accounts.
- An **Admin** has Member permissions and can invite other users.
- A **Member** can fully manage apps and clusters.
- A **Viewer** can view resources.

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 roles](#)

Copyright information

Copyright © 2024 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.