


SOLUTION BRIEF

# Simplify Jenkins application and data lifecycle management

Deliver data protection, disaster recovery, and portability to Jenkins



## Introduction

Many enterprises are moving from a monolithic software development approach to an agile one that encompasses continuous integration and continuous delivery (CI/CD). The CI/CD method provides for smaller changes to code very frequently. Immediately after small changes are checked in and merged, new builds and testing are completed on the new code via automation. This process allows developers to identify any potential issues easier and faster, as well as to get new updates to customers sooner. Jenkins is an open source tool that can be used to automate many different tasks related to building, testing, delivering, and/or deploying software. It also stores build metadata, testing results and may also store build artifacts.

However, when Jenkins is deployed on a Kubernetes cluster, the volumes holding the pipeline and other artifact data must always remain persistent. A new software build might not pass the automated tests, and developers might need to additional information regarding test results, so the all data must be continuously available. Further, the pipeline, with its build metadata and artifacts, may need to be relocated to a new cluster for portability and/or disaster recovery. Astra provides automated persistence and cloning for these scenarios.

## Astra overview

Astra is a fully managed service that makes it easier for our customers to manage, protect, and move their data-rich containerized workloads running on Kubernetes within and across public clouds and

## Key benefits headline

- Persist your Jenkins pipeline definitions and artifacts
- Recover all Jenkins data and Kubernetes resources in a disaster scenario
- Maintain your application binaries, pipeline configurations, and artifacts when you redeploy to a new Kubernetes cluster

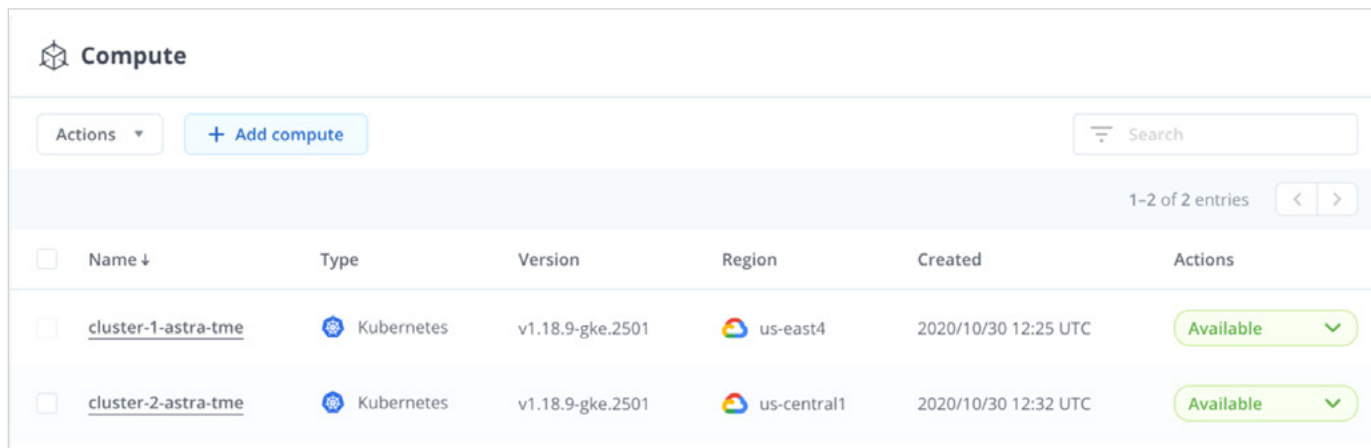
on premises. Astra provides persistent container storage that leverages NetApp's proven and expansive storage portfolio in the public cloud and on premises. It also offers a rich set of advanced application-aware data management functionality (like snapshot revert, backup and restore, activity log, and active cloning) for data protection, disaster recovery, data audit, and migration use cases for your modern apps.

## Managing Jenkins with Astra

Simply register your Google Kubernetes Engine (GKE) clusters in the cloud with Astra. Upon registration, Astra:

- Installs NetApp® Trident, NetApp's open source Kubernetes storage orchestrator
- Creates a bucket on the cloud object store for future backups
- Creates a service account on your cluster for itself

The following example shows two clusters, one located in the us-east4 GCP region (N. Virginia) and one located in the us-central1 (Iowa) region.



The screenshot shows the 'Compute' section of a cloud management console. It features a table with two entries for registered Kubernetes clusters. The table has columns for Name, Type, Version, Region, Created, and Actions. The first cluster is 'cluster-1-astra-tme' in the 'us-east4' region, created on 2020/10/30 at 12:25 UTC. The second cluster is 'cluster-2-astra-tme' in the 'us-central1' region, created on 2020/10/30 at 12:32 UTC. Both clusters are marked as 'Available'.

Name ↓	Type	Version	Region	Created	Actions
<a href="#">cluster-1-astra-tme</a>	Kubernetes	v1.18.9-gke.2501	us-east4	2020/10/30 12:25 UTC	Available
<a href="#">cluster-2-astra-tme</a>	Kubernetes	v1.18.9-gke.2501	us-central1	2020/10/30 12:32 UTC	Available

Figure 1: Registered clusters.

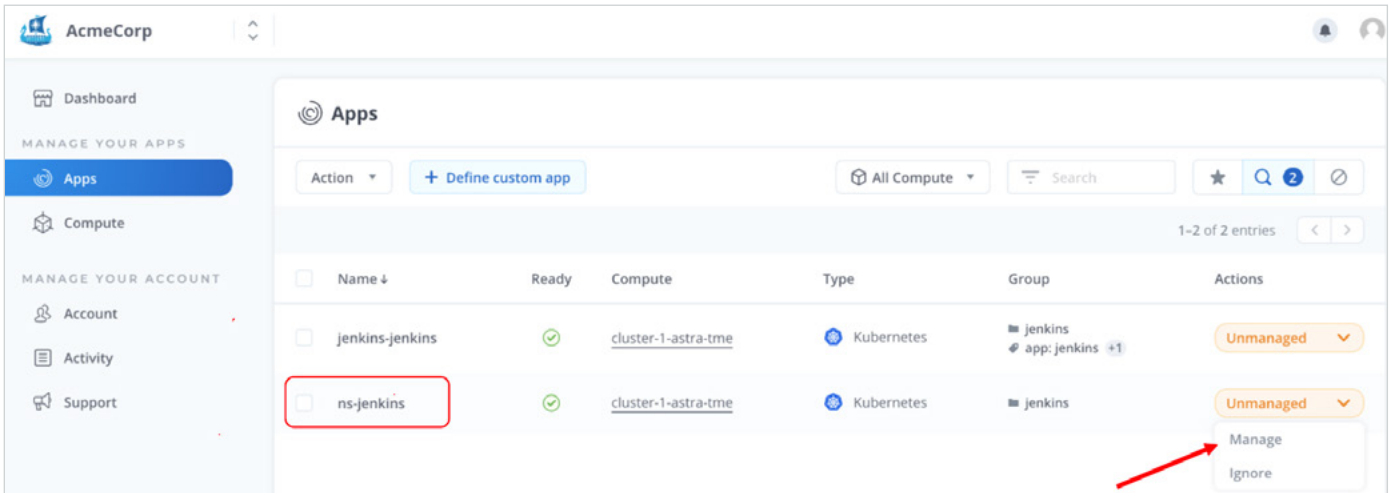


Figure 2: Managing the Jenkins namespace.

Install Jenkins on cluster-1-astra-tme (located at us-east4) using the current Bitnami Helm chart at bitnami/Jenkins or a custom manifest. (Support for the Jenkins Kubernetes plug-in is coming soon.) Astra discovers the applications on your registered clusters, and you can manage either just the application or all the resources in the entire namespace as one unit. Trident automatically generates the Kubernetes Persistent Volumes (PVs) and NetApp Cloud Volumes Service volumes that Jenkins needs.

Now Astra can take snapshots, backups, and clones of that application, its resources, and its associated PVs.

### Persisting Jenkins data

All the data saved by Jenkins can be automatically persisted by using snapshots and backups. Astra snapshots and backups preserve the application, its metadata, and its volumes in one easily manageable unit. The PV snapshot is stored in NetApp Cloud Volumes Service, and the application snapshot is stored in Astra. All PV backups are stored in the object store.

Both on-demand and scheduled snapshots and backups are supported. You can set up a snapshot and backup schedule for the volume and all the Kubernetes objects that are associated with Jenkins.

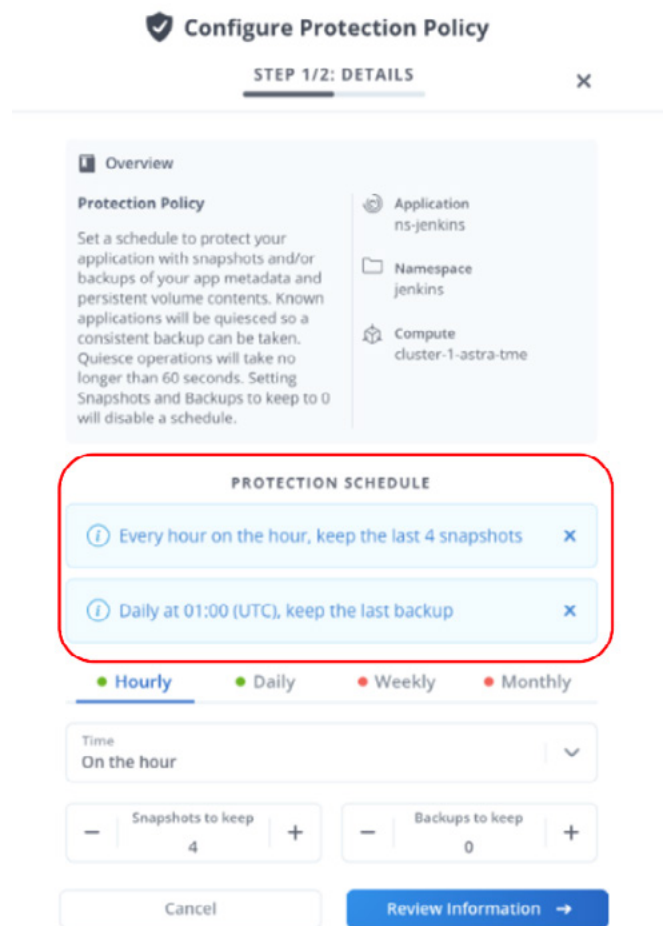


Figure 3: Configuring the protection policy.

After reviewing the information, set the protection policy. Astra automatically takes snapshots and backups based on the schedule you established. Of course, you can always take snapshots and backups on demand. You can even back up from a previous snapshot.

## Cloning Jenkins to a new cluster

Now you are protected from a disaster, such as losing the cluster or accidentally deleting the namespace. You also have the ability to redeploy Jenkins to a new namespace or a new cluster with the Jenkins state either as it was during the backup or in its current state.

For example, suppose that you have a project that will soon change ownership. The current team already has Jenkins configured with a link to the proper GitHub repository, artifacts, and a testing pipeline setup. They also have stored builds. You would like to redeploy Jenkins, along with all its data, closer to the new team in the Midwest. Jenkins is currently running on the cluster-1-astra-tme cluster in the us-east4 (N.Virginia) region and has three builds already made.

Next, clone Jenkins to an entirely new cluster, cluster-2-astra-tme, in a different Google Cloud Platform region using one of the backups. You could also clone at the current state, because the Astra clone application first creates a backup for you if you don't specify one. This cloning brings up a new instance of Jenkins, running at the same state as when the backup was taken.

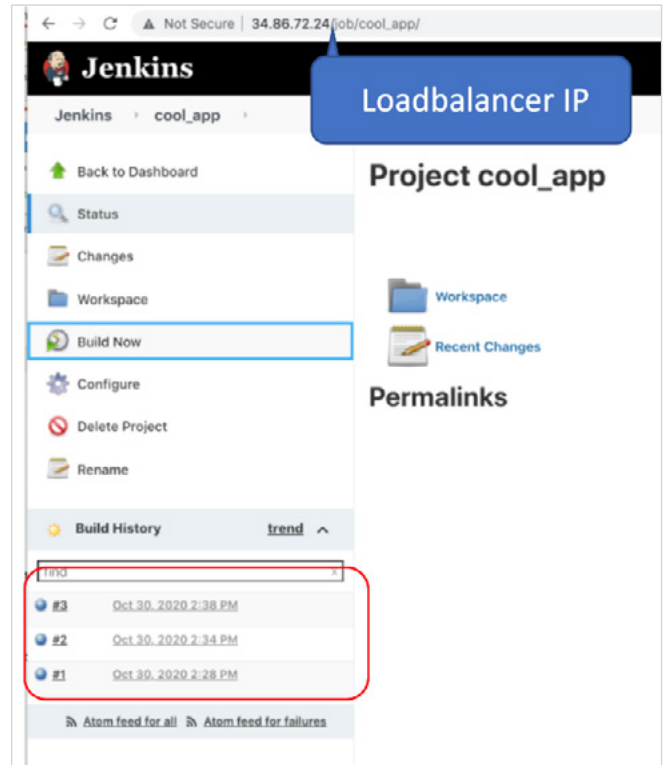


Figure 4: Current Jenkins state on N. Virginia cluster.

The new Jenkins clone will be provisioned automatically, managed by Astra, and ready with all its persisted data from the original instance.

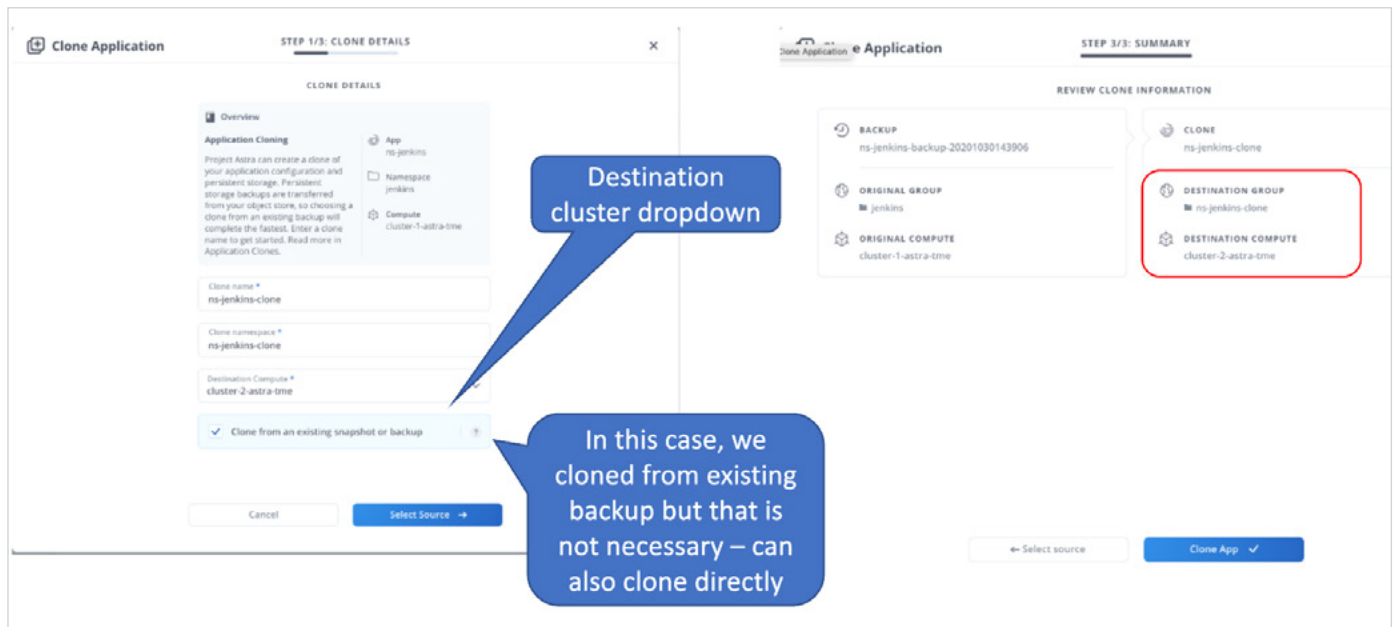


Figure 5: Cloning Jenkins using a backup.

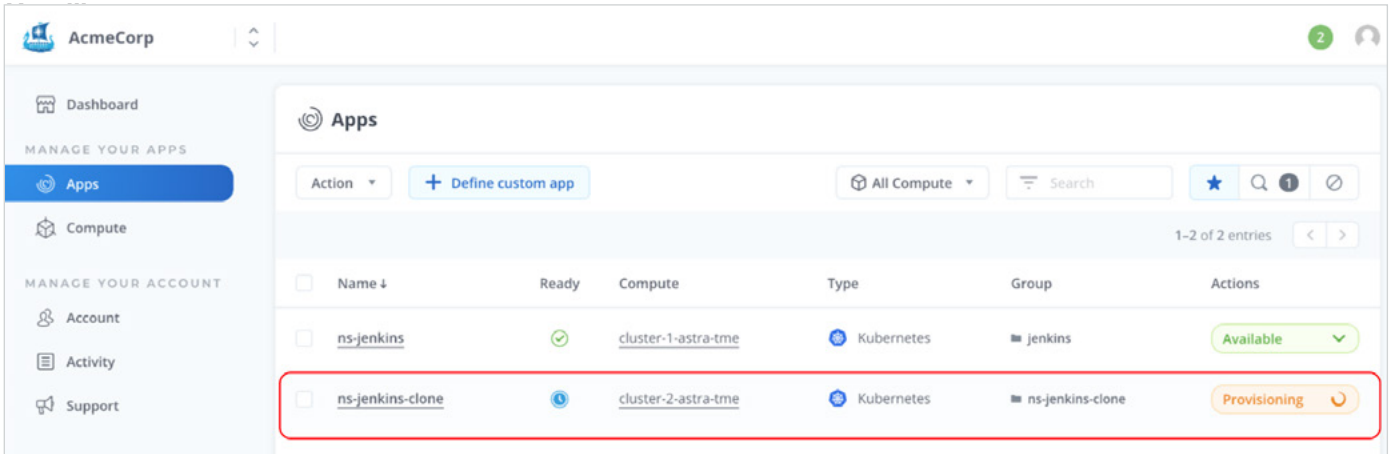


Figure 6: Provisioning the new Jenkins instance on a new cluster.

When the new Jenkins instance comes up, you will have the current build metadata that was there with the cloned instance, and you can run a new build without having to reconfigure the Jenkins project from the beginning.

### How can I learn more?

To learn more, visit the [Astra website](#) and the [documentation](#) on Astra.

### About NetApp

In a world full of generalists, NetApp is a specialist. We're focused on one thing, helping your business get the most out of your data. NetApp brings the enterprise-grade data services you rely on into the cloud, and the simple flexibility of cloud into the data center. Our industry-leading solutions work across diverse customer environments and the world's biggest public clouds.

As a cloud-led, data-centric software company, only NetApp can help build your unique data fabric, simplify and connect your cloud, and securely deliver the right data, services and applications to the right people—anytime, anywhere.

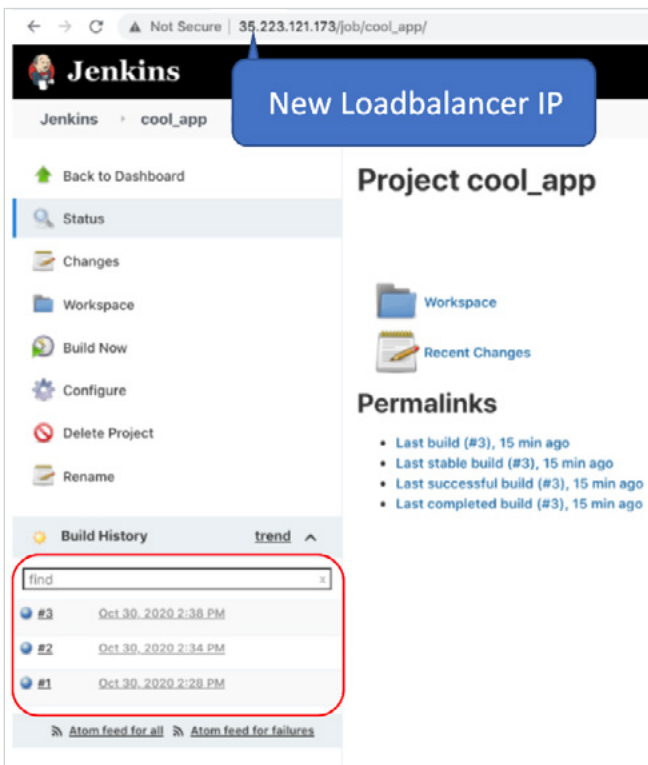


Figure 7: New Jenkins instance in a new cluster.

