


SOLUTION BRIEF

Simplify Jenkins application and data lifecycle management

Deliver data protection, restoration, and portability to Jenkins on Google Kubernetes Engine



Key benefits

- Persist your Jenkins pipeline definitions and artifacts
- Restore all Jenkins data and Kubernetes resources in a disaster scenario
- Maintain your build data, pipeline configurations, and artifacts when you relocate to a new Kubernetes cluster or namespace

Introduction

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.

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 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 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 and three Kubernetes storage classes using NetApp Cloud Volumes Service.
- Creates a bucket on the cloud object store for future backups for all your registered clusters
- Creates a service account on your cluster for itself

As an example, Figure 1 shows two clusters registered with Astra, one located in the europe-west2 Google Cloud Platform (GCP) region (London) and one located in the us-west2 (Los Angeles) GCP region.

Clone Jenkins to a new namespace or cluster along with all its current pipeline definitions, build data, and artifacts.

The screenshot shows the 'Compute' section of a management console. It features a table with columns for Name, Type, Version, Location, Created, and Actions. Two clusters are listed: 'cluster-1-patd' and 'cluster-2-patd', both using Kubernetes v1.19.6-gke.600. The first is located in europe-west2-a and the second in us-west2-b. Both are in an 'Available' state.

Name	Type	Version	Location	Created	Actions
cluster-1-patd	Kubernetes	v1.19.6-gke.600	europe-west2-a	2021/01/21 15:23 UTC	Available
cluster-2-patd	Kubernetes	v1.19.6-gke.600	us-west2-b	2021/01/21 15:24 UTC	Available

Figure 1) Registered clusters.

After your cluster is registered, install Jenkins on your cluster using the current Bitnami Helm chart at bitnami/Jenkins or a custom manifest. (Support for the Jenkins Kubernetes plug-in is coming soon. [See Cloning Jenkins Kubernetes plugin with Astra](#)) Trident automatically generates the Kubernetes Persistent Volumes (PVs) and NetApp Cloud Volumes Service (CVS) volumes that Jenkins needs.

Astra discovers all the applications on your registered clusters. You can manage together just the application, all the resources in the entire namespace, or a custom group based on object labels. Figure 2 shows managing the entire Jenkins namespace as one unit.

The screenshot shows the 'Apps' section of the management console. A table lists discovered applications. The 'jenkins' application is highlighted with a red box. Its 'Actions' dropdown menu is open, showing options for 'Unmanaged', 'Manage', and 'Ignore'. A red arrow points to the 'Manage' option.

Name	Ready	Compute	Type	Group	Discovered	Actions
jenkins	✓	cluster-1-patd	Kubernetes	jenkins	2021/01/21 15:45 UTC	Unmanaged Manage Ignore
jenkins-jenkins	✓	cluster-1-patd	Kubernetes	jenkins app.kubernetes.io/name: jenkins +1	2021/01/21 15:45 UTC	

Figure 2) Managing the Jenkins namespace.

Log in, configure, and use Jenkins as you normally would; you can create builds, configure pipelines, etc.

Persisting Jenkins data

All Jenkins data 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 Volume Services (CVS), 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, as shown in Figure 3.

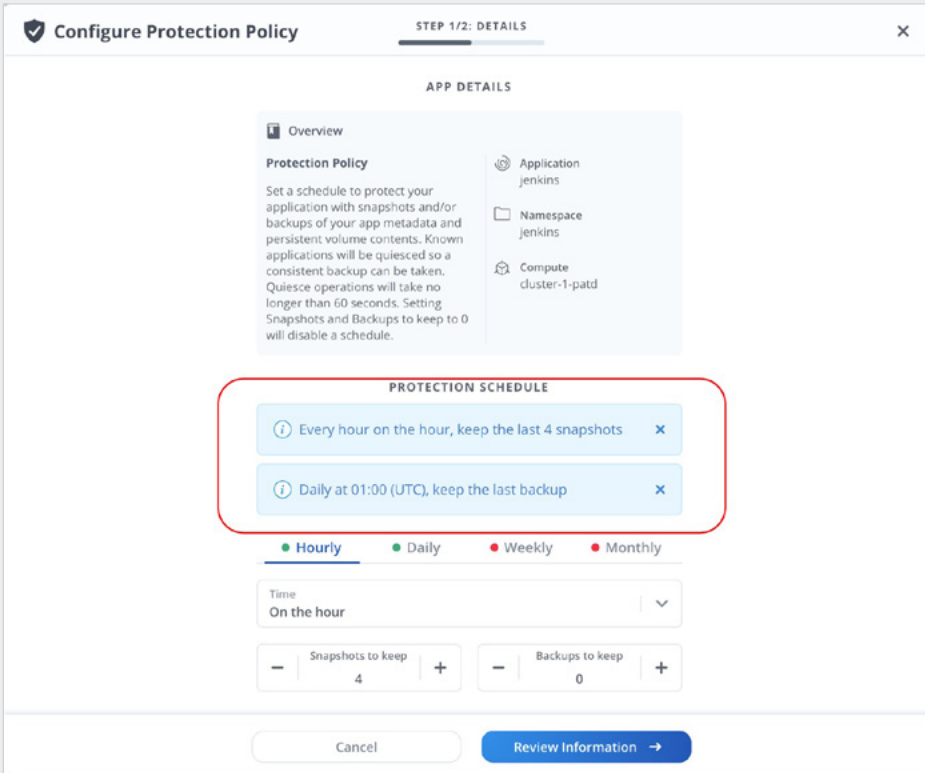


Figure 3) Configuring the protection policy.

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 since you can restore from backup to a new cluster or namespace. You also have the ability to redeploy Jenkins to a new namespace or a new cluster at the current Jenkins state with a direct clone.

For example, suppose that you have a project that will soon change ownership. The current team in London already has Jenkins configured with a link to the proper GitHub repository, artifacts, and a testing pipeline setup. They also have stored build metadata. You would like to redeploy Jenkins, along with all its data, closer to the new team in Los Angeles. Jenkins is currently running on the cluster-1-patd cluster in the europe-west2 (London) region and has three builds already made as shown in Figure 4.

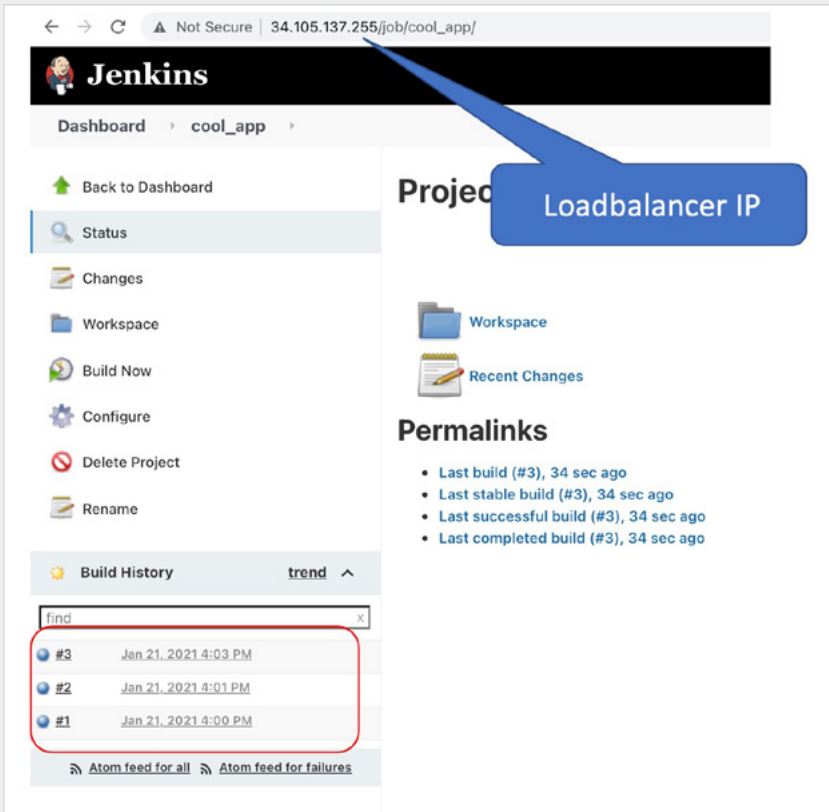


Figure 4 – Current Jenkins state on the London cluster.

Ensure a new cluster is registered with Astra. Our example in Figure 1 shows cluster-2-patd registered in GCP us-west2 (Los Angeles) region.

Next, use the drop down menu to clone Jenkins and its resources as depicted in Figure 5

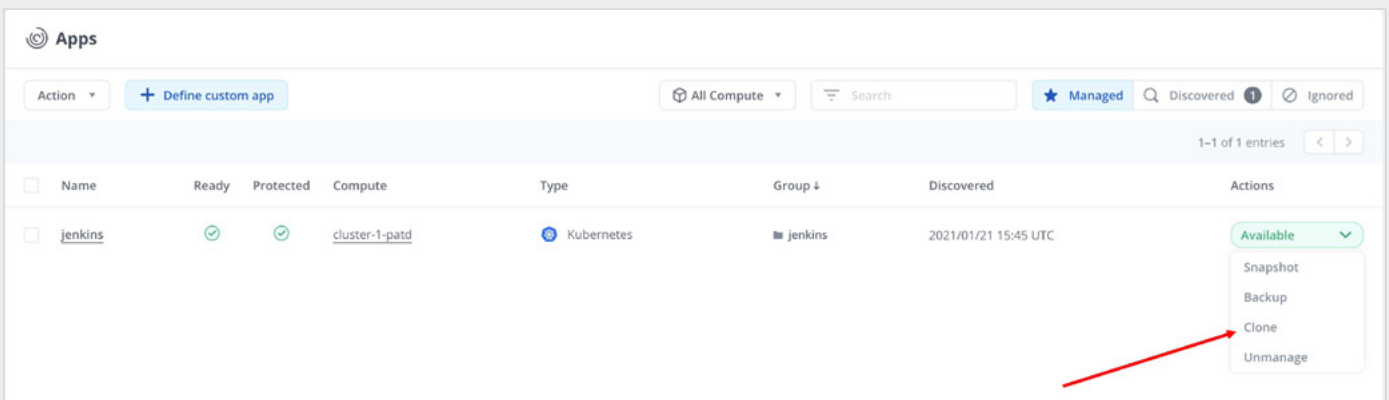


Figure 5) Clone Jenkins and its resources.

Clone Jenkins to the new cluster, cluster-2-patd, using a backup as shown in Figure 6. Cloning from backup brings up a new instance of Jenkins on the new cluster, running at the same state as when the backup was taken.

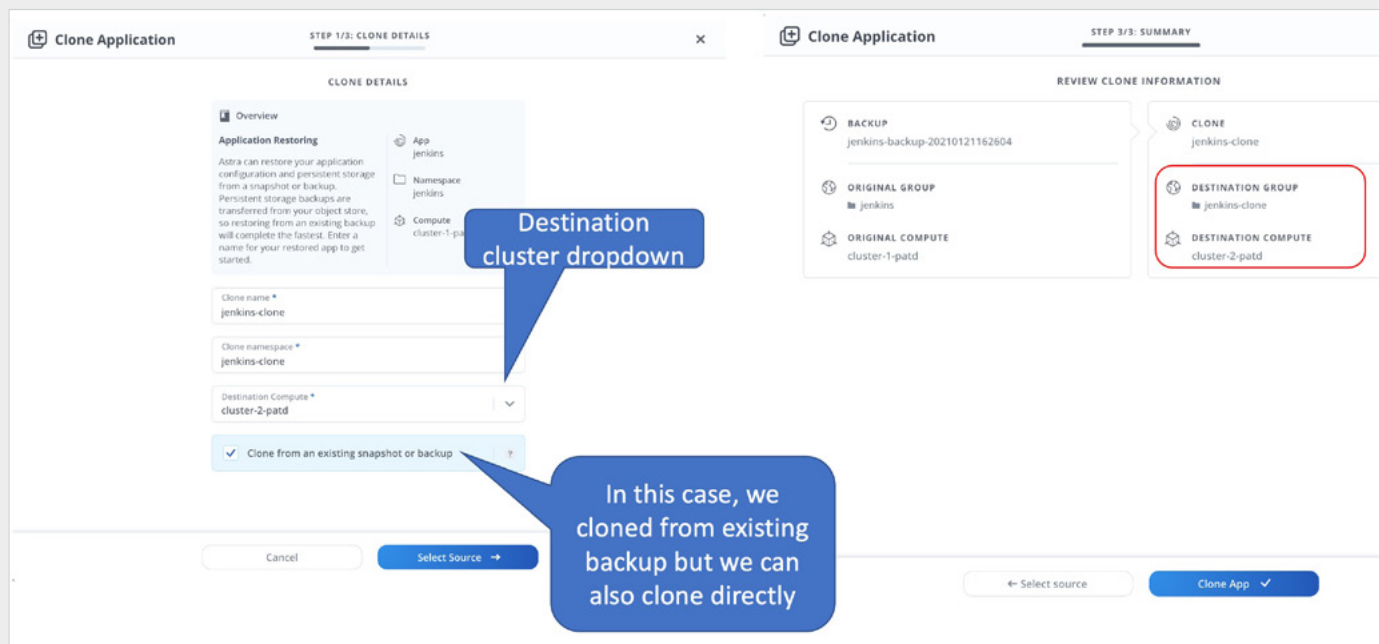


Figure 6) Cloning Jenkins using a backup.

You could also clone at the current state, which integrates the backup and cloning together.

The new Jenkins clone on the new cluster will be provisioned and managed automatically by Astra, as shown in Figure 7.

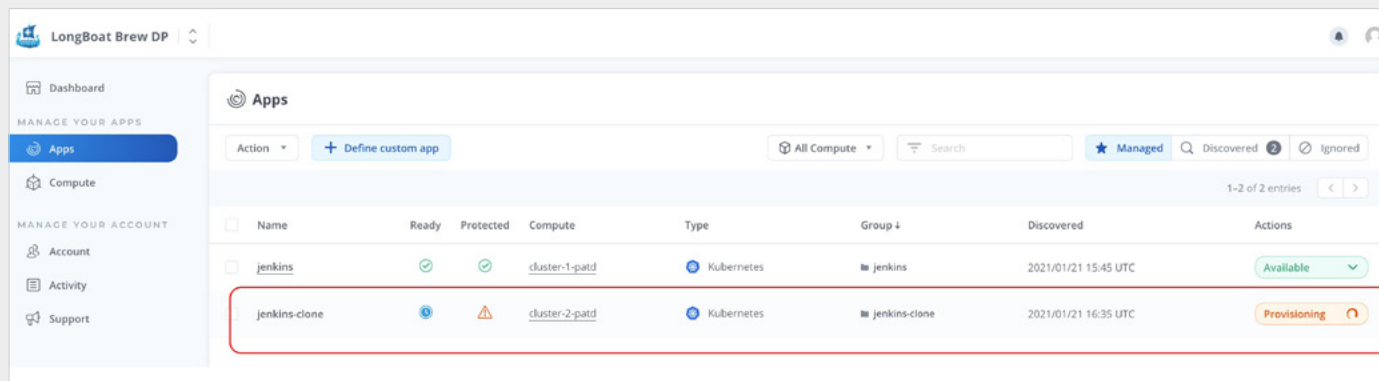


Figure 7) Provisioning the new Jenkins instance on a new cluster.

The new Jenkins instance will have all the data that was there with the cloned instance as shown in Figure 8. You can run a new build on the clone without having to reconfigure the Jenkins project from the beginning.

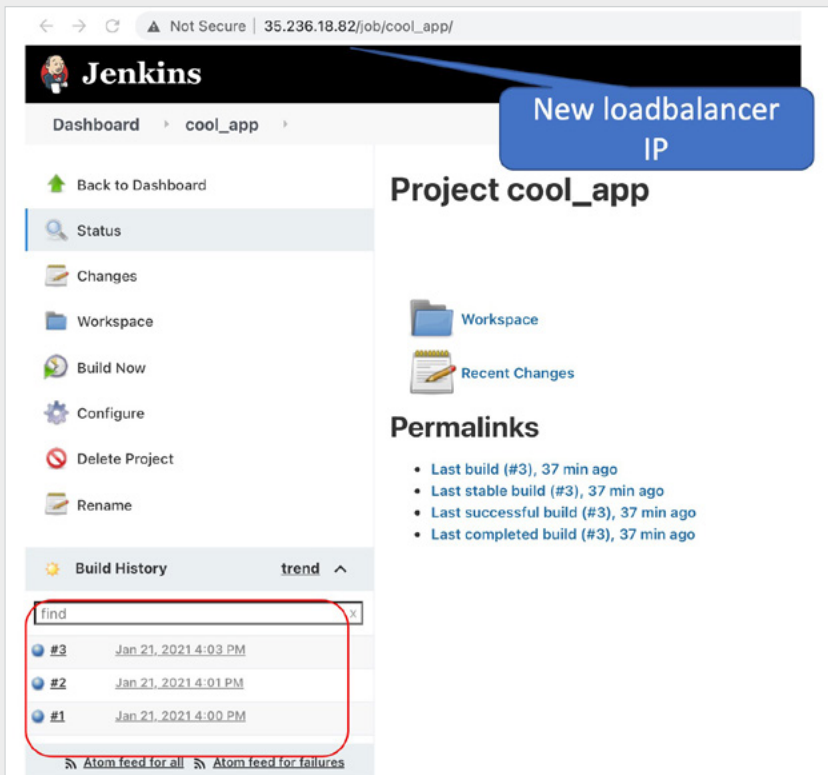



Figure 8) New Jenkins instance in a new cluster.

Where can I learn more?

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

