



Solution validation

FlexPod

NetApp
October 30, 2025

Table of Contents

Solution validation	1
Overview	1
Application recovery with remote backups.....	1
Data protection with Snapshot copies and application mobility for DevTest	13

Solution validation

Overview

[Previous: Astra Control Center installation on OpenShift Container Platform.](#)

In this section, we revisit the solution with some use cases:

- Restoring a stateful application from a remote backup to another OpenShift cluster running in the cloud.
- Restoring a stateful application to the same namespace in the OpenShift cluster.
- Application mobility by cloning from one FlexPod system (OpenShift Container Platform Bare Metal) to another FlexPod system (OpenShift Container Platform on VMware).

Notably, only a few use cases are validated in this solution. This validation does not in any way represent the entire functionality of Astra Control Center.

[Next: Application recovery with remote backups.](#)

Application recovery with remote backups

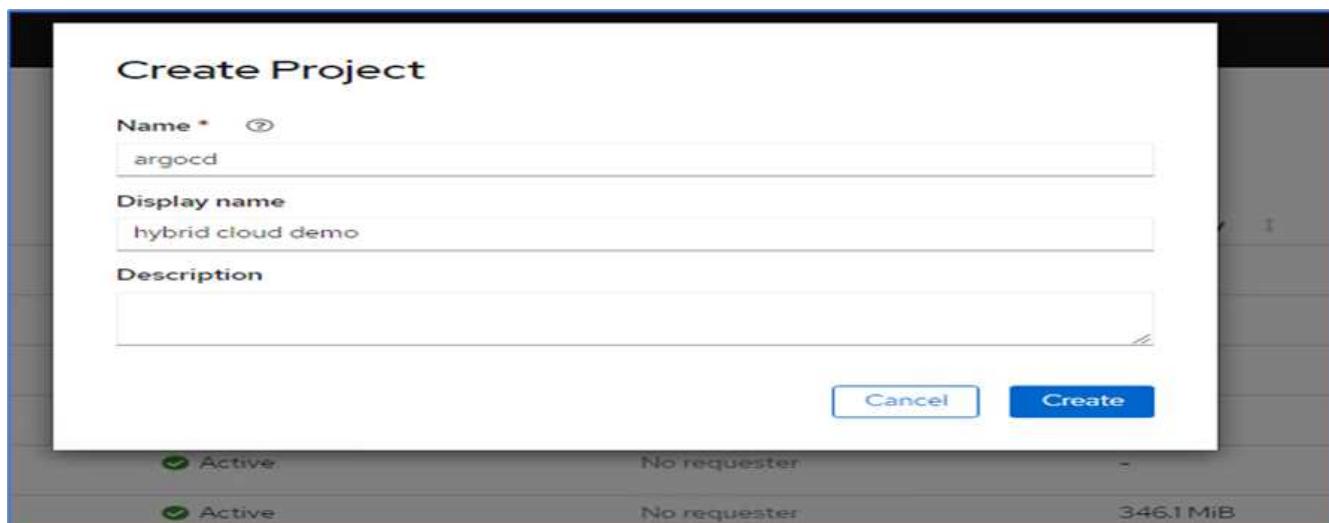
[Previous: Solution validation overview.](#)

With Astra, you can take a full application-consistent backup that can be used to restore your application with its data to a different Kubernetes cluster running in an on-premises data center or in a public cloud.

To validate a successful application recovery, simulate an on-premises failure of an application running on the FlexPod system and restore the application to a K8s cluster running in the cloud by using a remote backup.

The sample application is a pricelist application that uses MySQL for the database. To automate the deployment, we used the [Argo CD](#) tool. Argo CD is a declarative, GitOps, continuous delivery tool for Kubernetes.

1. Log into the on-premises OpenShift cluster and create a new project with the name `argocd`.



2. In the OperatorHub, search for argocd and select Argo CD operator.

Project: argocd

OperatorHub

Discover Operators from the Kubernetes community and Red Hat partners, curated by Red Hat. You can purchase commercial software installation, the Operator capabilities will appear in the Developer Catalog providing a self-service experience.

All Items

argocd

 Argo CD provided by Argo CD Community	 Argo CD Operator (Helm) provided by Disposable Zone
Argo CD is a declarative, GitOps continuous delivery tool for Kubernetes.	Declarative Continuous Delivery following Gitops.

All Items

argocd

3. Install the operator in the argocd namespace.

OperatorHub > Operator Installation

Install Operator

Install your Operator by subscribing to one of the update channels to keep the Operator up to date. The strategy determines either manual or automatic updates.

Update channel * ⓘ

alpha

Installation mode *

All namespaces on the cluster (default)
Operator will be available in all Namespaces.

A specific namespace on the cluster
Operator will be available in a single Namespace only.

Installed Namespace *

argocd

Update approval * ⓘ

Automatic

Manual

Argo CD
provided by Argo CD Community

Provided APIs

Application An Application is a group of Kubernetes resources as defined by a manifest.	ApplicationSet An ApplicationSet is a group or set of Application resources.
AppProject An AppProject is a logical grouping of Argo CD Applications.	ArgoCDE Argo CDE ArgoCDE is the Schema for the argocdexports API
Argo CD ArgoCD is the Schema for the argocds API	

Install **Cancel**

4. Go to the operator and click Create ArgoCD.

Project: argocd

Installed Operators > Operator details

 Argo CD
0.3.0 provided by Argo CD Community

Actions ▾

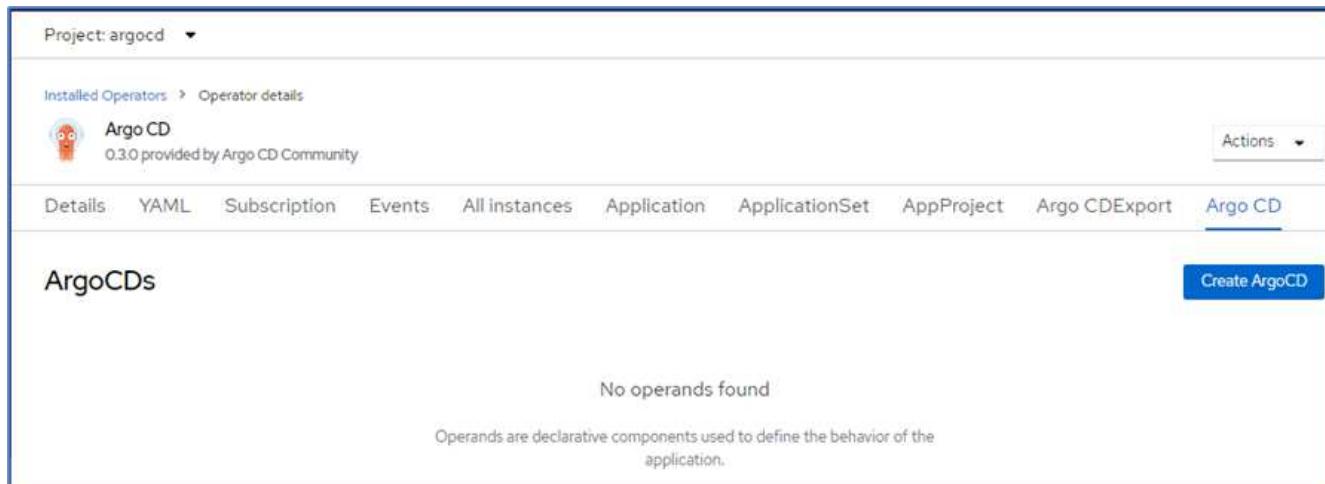
Details YAML Subscription Events All instances Application ApplicationSet AppProject Argo CDExport Argo CD

ArgoCDs

Create ArgoCD

No operands found

Operands are declarative components used to define the behavior of the application.



5. To deploy the Argo CD instance in the `argocd` project, provide a name and click Create.

Project: argocd

ArgoCD > Create ArgoCD

Create ArgoCD

Create by completing the form. Default values may be provided by the Operator authors.

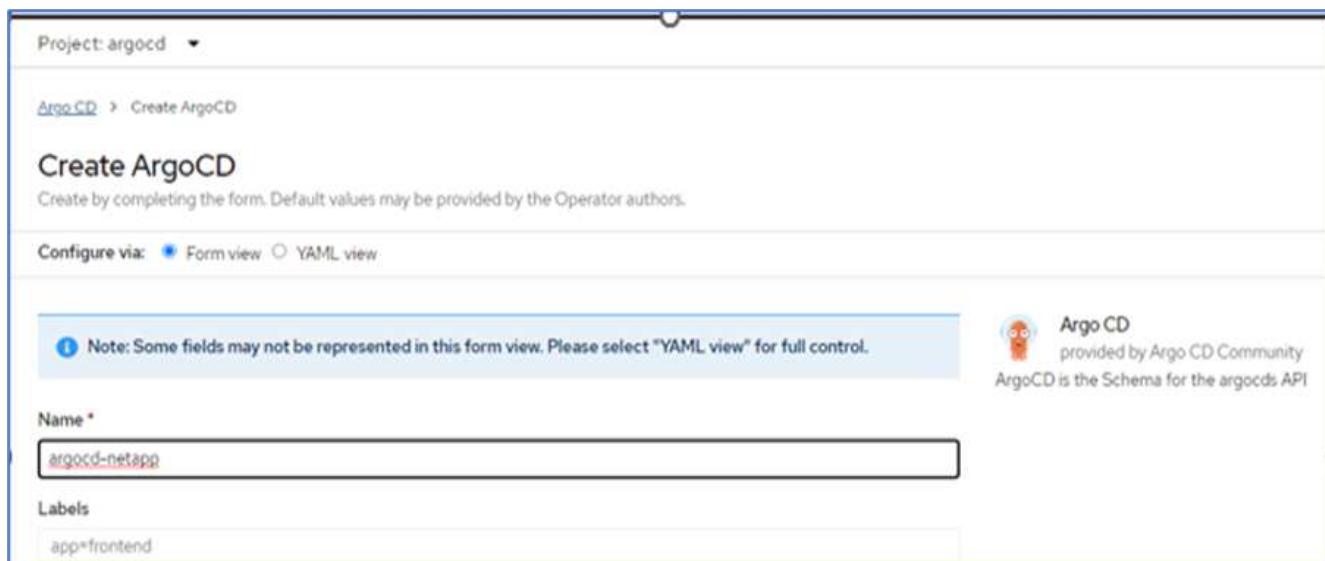
Configure via: Form view YAML view

Note: Some fields may not be represented in this form view. Please select "YAML view" for full control.

Name *

Labels

 Argo CD
provided by Argo CD Community
ArgoCD is the Schema for the argocds API



6. To log in to Argo CD, the default user is admin and the password is in a secret file with the name `argocd-netapp-cluster`.

Project: argocd

Secrets > Secret details

argocd-netapp-cluster
Managed by  argocd-netapp

Details YAML

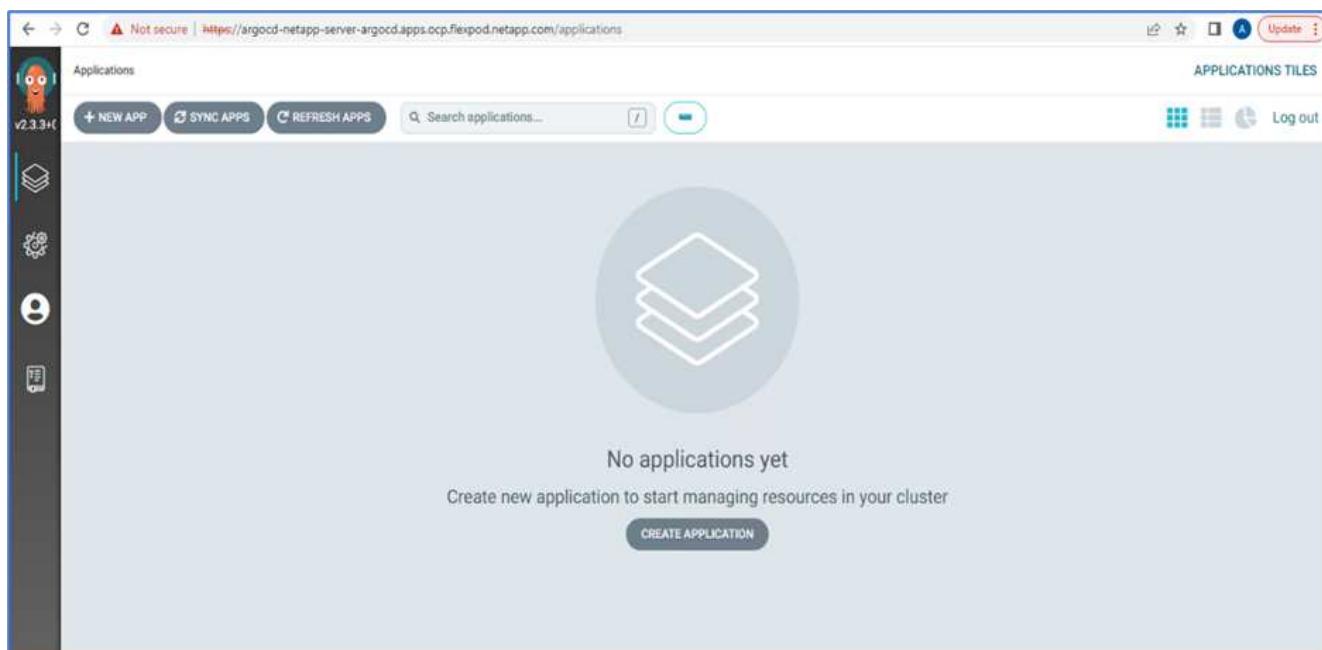
Secret details

Name	argocd-netapp-cluster	Type	Opaque
Namespace	NS argocd		
Labels	app.kubernetes.io/managed-by=argocd-netapp, app.kubernetes.io/name=argocd-netapp-cluster, app.kubernetes.io/part-of=argocd	Edit	
Annotations	argocd.argocd.netapp		
Created at	2 minutes ago		
Owner	 argocd-netapp		

Data

admin.password	*****	 Reveal values	 Copied	
----------------	-------	---	--	---

7. From the side menu, select Routes > Location and click the URL for the `argocd` routes. Enter the user name and password.



Not secure | <https://argocd-netapp-server-argocd.apps.ocp.flexpod.netapp.com/applications>

Applications

APPLICATIONS TILES

v2.3.3-h

+ NEW APP  SYNC APPS  REFRESH APPS

Search applications...

No applications yet

Create new application to start managing resources in your cluster

CREATE APPLICATION

Log out

8. Add the on-premises OpenShift cluster to Argo CD through the CLI.

```

#####Login to Argo CD#####
abhinav3@abhinav-ansible$ argocd-linux-amd64 login argocd-netapp-server-
argocd.apps.ocp.flexpod.netapp.com --insecure
Username: admin
Password:
'admin:login' logged in successfully
Context 'argocd-netapp-server-argocd.apps.ocp.flexpod.netapp.com' updated
#####List the On-Premises OpenShift cluster#####
abhinav3@abhinav-ansible$ argocd-linux-amd64 cluster add
ERRO[0000] Choose a context name from:
CURRENT NAME
CLUSTER SERVER
* default/api-ocp-flexpod-netapp-com:6443/abhinav3
api-ocp-flexpod-netapp-com:6443
https://api.ocp.flexpod.netapp.com:6443
      default/api-ocp1-flexpod-netapp-com:6443/abhinav3
api-ocp1-flexpod-netapp-com:6443
https://api.ocp1.flexpod.netapp.com:6443
#####Add On-Premises OpenShift cluster#####
abhinav3@abhinav-ansible$ argocd-linux-amd64 cluster add default/api-
ocp1-flexpod-netapp-com:6443/abhinav3
WARNING: This will create a service account `argocd-manager` on the
cluster referenced by context `default/api-ocp1-flexpod-netapp-
com:6443/abhinav3` with full cluster level admin privileges. Do you want
to continue [y/N]? y
INFO[0002] ServiceAccount "argocd-manager" already exists in namespace
"kube-system"
INFO[0002] ClusterRole "argocd-manager-role" updated
INFO[0002] ClusterRoleBinding "argocd-manager-role-binding" updated
Cluster 'https://api.ocp1.flexpod.netapp.com:6443' added

```

9. In the ArgoCD UI, click NEW APP and enter the details about the app name and code repository.

CREATE CANCEL X EDIT AS YAML

GENERAL

Application Name: pricelist

Project: default

SYNC POLICY

Manual

SYNC OPTIONS

SKIP SCHEMA VALIDATION AUTO-CREATE NAMESPACE

PRUNE LAST APPLY OUT OF SYNC ONLY

RESPECT IGNORE DIFFERENCES

PRUNE PROPAGATION POLICY: foreground

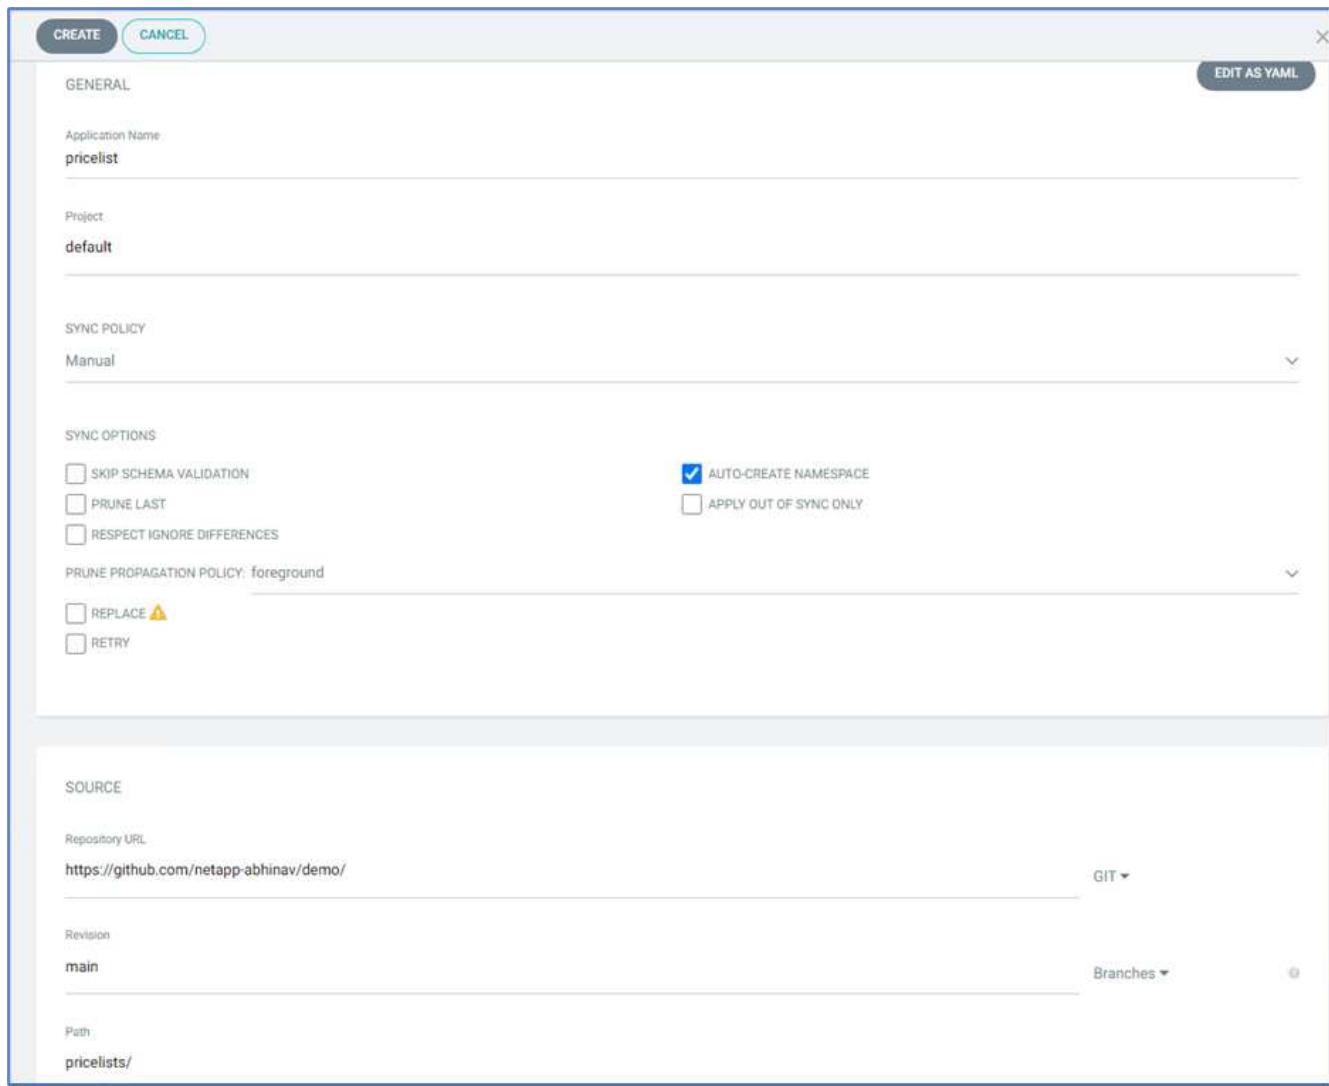
REPLACE ⚠ RETRY

SOURCE

Repository URL: <https://github.com/netapp-abhinav/demo/> GIT

Revision: main Branches 0

Path: pricelists/

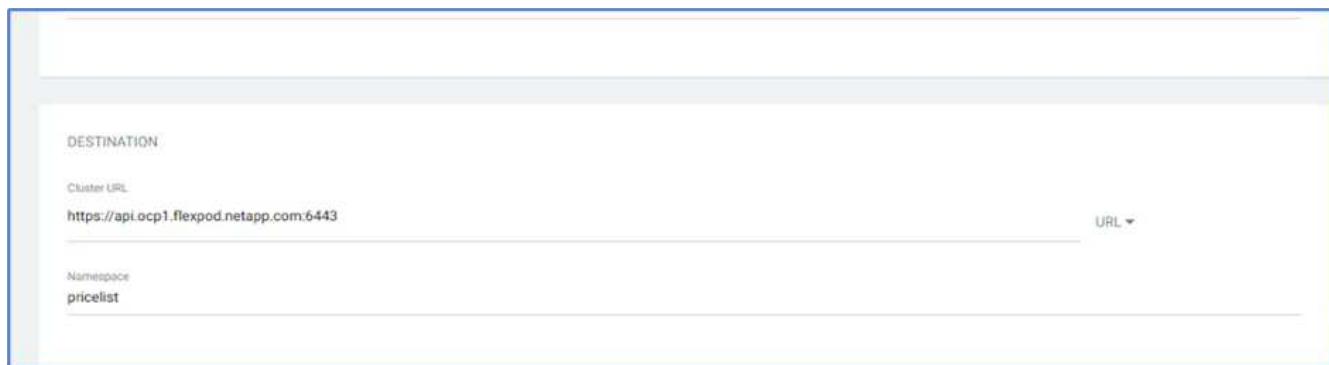


10. Enter the OpenShift cluster where the app will be deployed along with the namespace.

DESTINATION

Cluster URL: <https://api.ocp1.flexpod.netapp.com:6443> URL

Namespace: pricelist



11. To deploy the app on the on-premises OpenShift cluster, click SYNC.

12. In the OpenShift Container Platform console, go to Project Pricelist, and, under Storage, verify the name and size of the PVC.

13. Log into System Manager and verify the PVC.

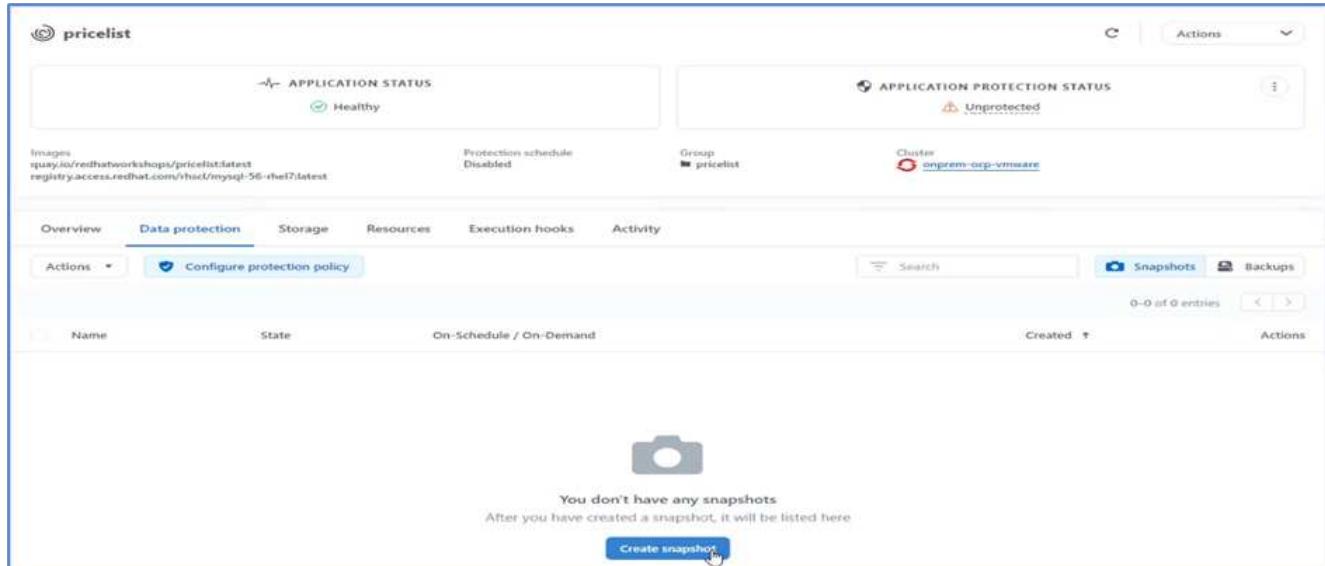
14. After the Pods are running, select Networking > Routes from the side menu, and click the URL under Location.

15. The Pricelist app homepage is displayed.

16. Create a few records on the web page.

17. The app is discovered in Astra Control Center. To manage the app, go to Applications > Discovered, select the Pricelist app, and click Manage Applications under Actions.

18. Click the Pricelist app and select Data Protection. At this point, there should be no snapshots or backups. Click Create Snapshot to create an on-demand snapshot.



APPLICATION STATUS
Healthy

APPLICATION PROTECTION STATUS
Unprotected

Images: quay.io/redhatworkshop/pricelist:latest
registry.access.redhat.com/rhsc1/mysql-56-rhel7:latest

Protection schedule: Disabled
Group: pricelist
Cluster: onprem-ocp-vmware

Overview Data protection Storage Resources Execution hooks Activity

Actions Configure protection policy

Search: Snapshots Backups

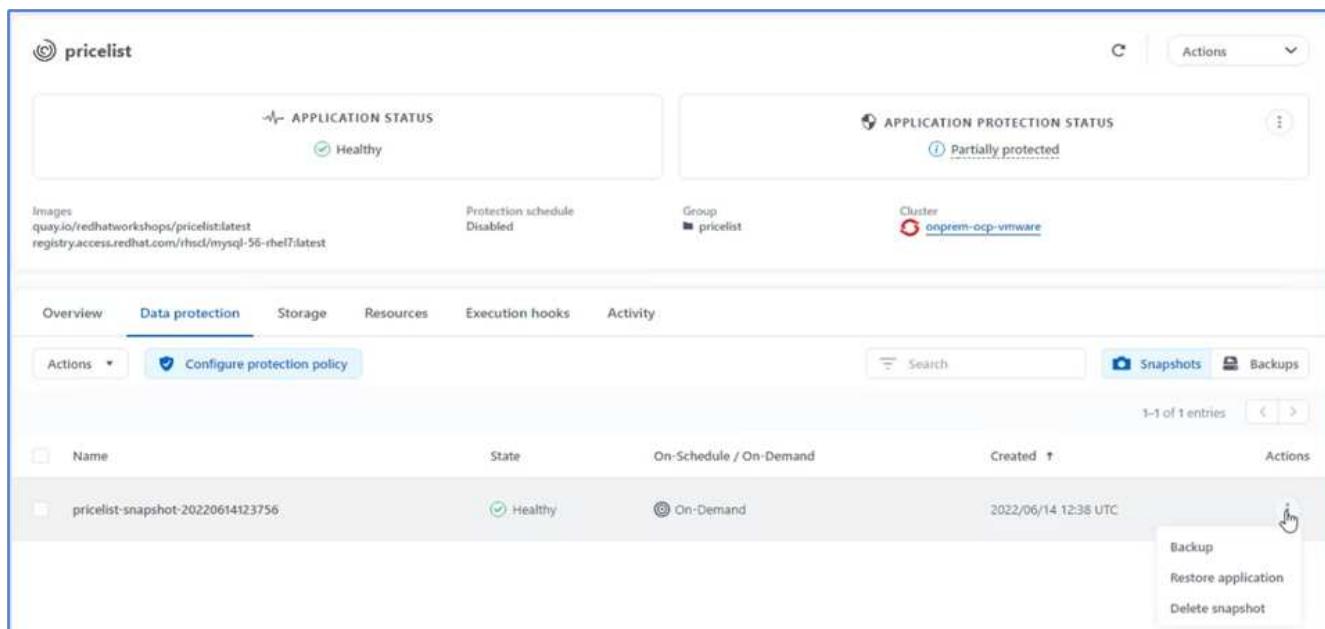
0-0 of 0 entries

Name	State	On-Schedule / On-Demand	Created	Actions
 You don't have any snapshots After you have created a snapshot, it will be listed here Create snapshot				



NetApp Astra Control Center supports both on-demand and scheduled snapshots and backups.

19. After the snapshot is created and the State is healthy, create a remote backup using that snapshot. This backup is stored in the S3 bucket.



APPLICATION STATUS
Healthy

APPLICATION PROTECTION STATUS
Partially protected

Images: quay.io/redhatworkshop/pricelist:latest
registry.access.redhat.com/rhsc1/mysql-56-rhel7:latest

Protection schedule: Disabled
Group: pricelist
Cluster: onprem-ocp-vmware

Overview Data protection Storage Resources Execution hooks Activity

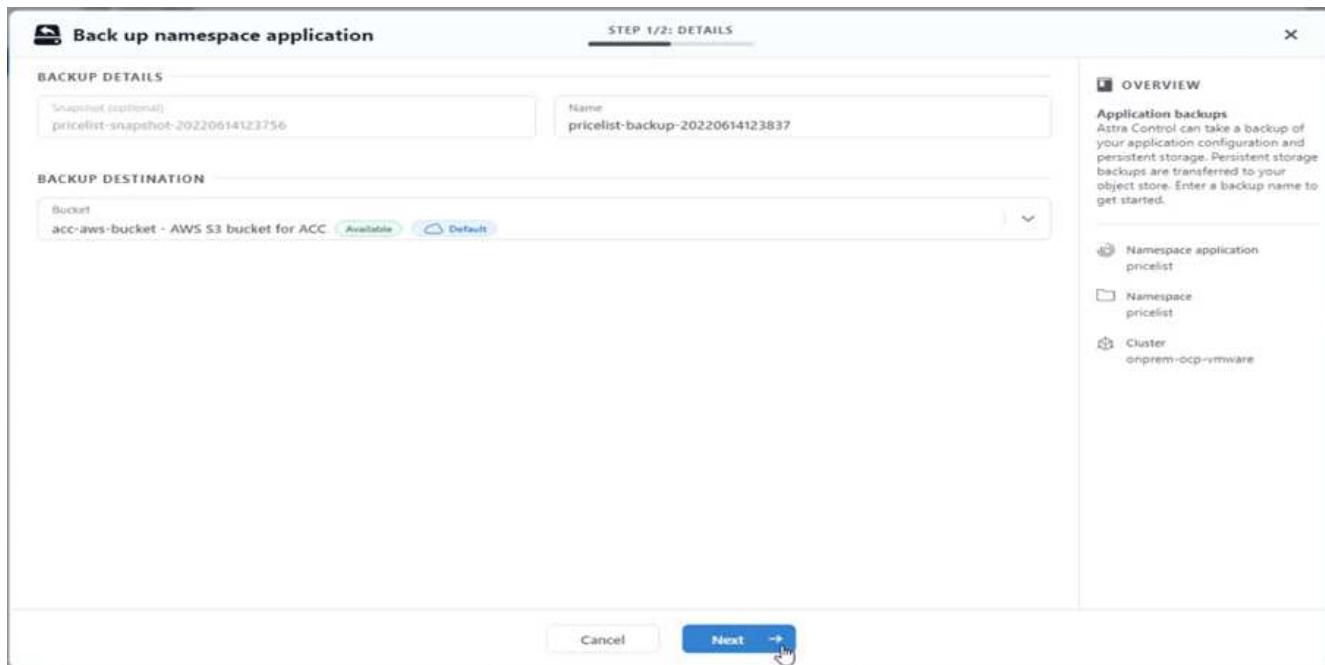
Actions Configure protection policy

Search: Snapshots Backups

1-1 of 1 entries

Name	State	On-Schedule / On-Demand	Created	Actions
pricelist-snapshot-20220614123756	Healthy	On-Demand	2022/06/14 12:38 UTC	Backup Restore application Delete snapshot

20. Select the AWS S3 bucket and initiate the backup operation.



21. The backup operation should create a folder with multiple objects in the AWS S3 bucket.

Name	Type	Last modified	Size	Storage class
config	-	June 14, 2022, 05:39:19 (UTC-07:00)	155.0 B	Standard
data/	Folder	-	-	-
index/	Folder	-	-	-
keys/	Folder	-	-	-
snapshots/	Folder	-	-	-

22. When the remote backup is complete, simulate a disaster on the on-premises by stopping the storage virtual machine (SVM) that hosts the backing volume for the PV.

Name	State	Subtype	Configured Protocols	IPspace
Infra_SVM	stopped	default	Default	Default

23. Refresh the webpage to confirm the outage. The webpage is unavailable.



As expected, the website is down, so let's quickly recover the app from the remote backup by using Astra to the OpenShift cluster running in AWS.

24. In Astra Control Center, click the Pricelist app and select Data Protection > Backups. Select the backup, and click Restore Application under Action.

Name	State	On-Schedule / On-Demand	Bucket	Created	Actions
pricelist-backup-20220614123837	Healthy	On-Demand	acc-aws-bucket	2022/06/14 12:38 UTC	Restore application Delete backup

25. Select ocp-aws as the destination cluster and give a name to the namespace. Click the on-demand backup, Next, and then Restore.

26. A new app with the name `pricelist-app` is provisioned on the OpenShift cluster running in AWS.

27. Verify the same in the OpenShift web console.

28. After all the pods under the `pricelist-aws` project are running, go to Routes and click the URL to launch the web page.

This process validates that the pricelist application has been successfully restored and that data integrity has been maintained on the OpenShift cluster running seamlessly on AWS with the help of Astra Control Center.

Data protection with Snapshot copies and application mobility for DevTest

This use case consists of two parts, as described the following sections.

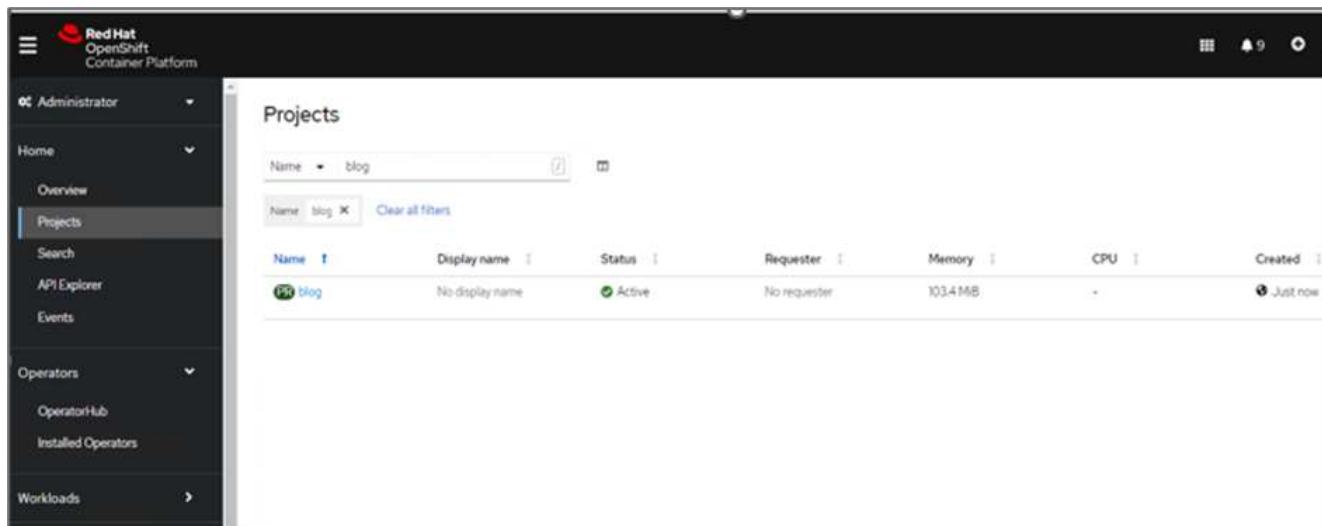
Part 1

With Astra Control Center, you can take application-aware snapshots for local data protection. If you accidentally delete or corrupt your data, you can revert your applications and associated data to a known good state using a previously recorded snapshot.

In this scenario, a development and testing (DevTest) team deploys a sample stateful application (blog site) that is a Ghost blog application, adds some content, and upgrades the app to the latest version available. The Ghost application uses SQLite for the database. Before upgrading the application, a snapshot (on-demand) is taken using Astra Control Center for data protection. The detailed steps are as follows:

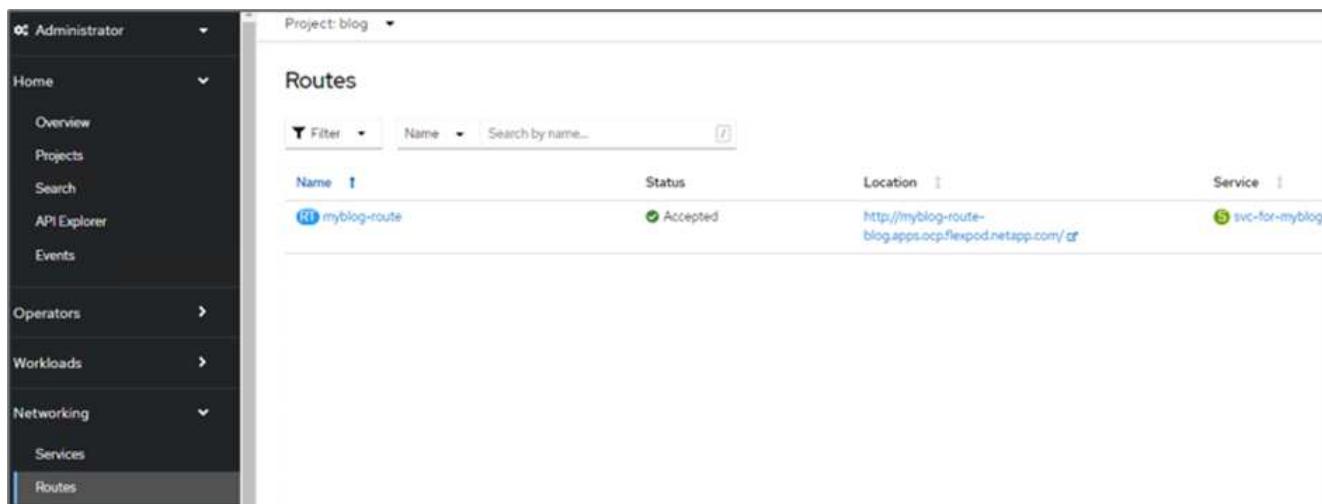
1. Deploy the sample blogging app and sync it from ArgoCD.

2. Log into the first OpenShift cluster, go to Project, and enter Blog in the search bar.



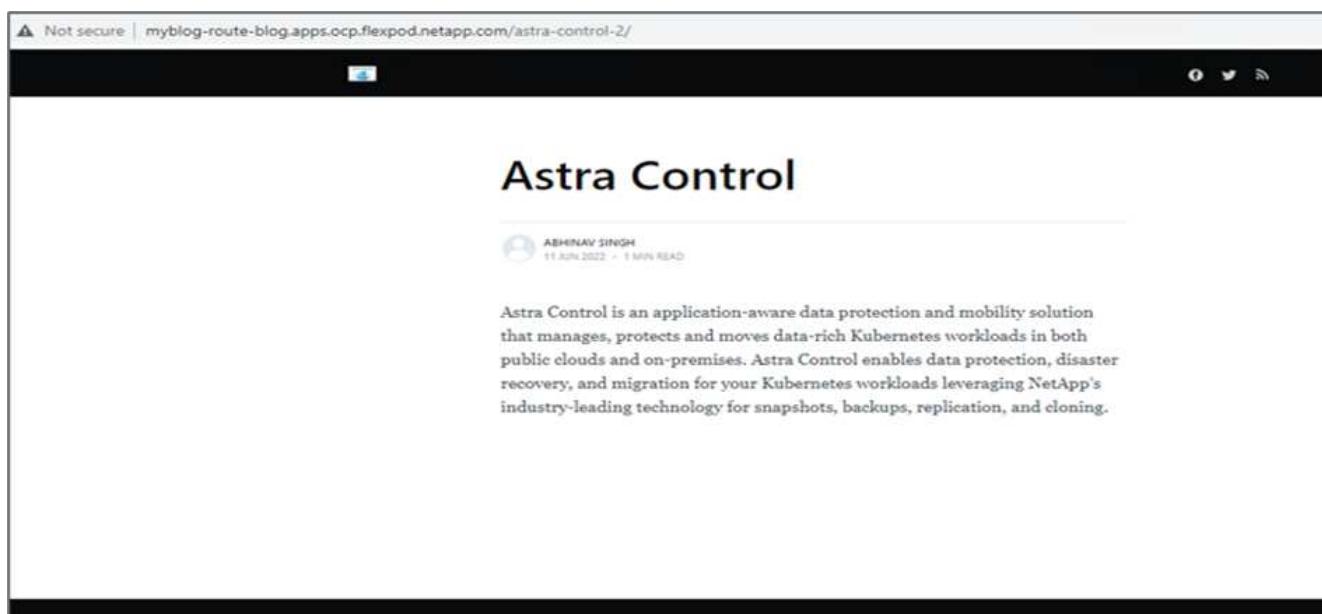
Name	Display name	Status	Requester	Memory	CPU	Created
blog	No display name	Active	No requester	103.4 MB	0	Just now

3. From the side menu, select Networking > Routes and click the URL.



Name	Status	Location	Service
myblog-route	Accepted	http://myblog-route-blog.apps.ocp.flexpod.netapp.com/or	svc-for-myblog

4. The blog home page is displayed. Add some content to the blog site and publish it.

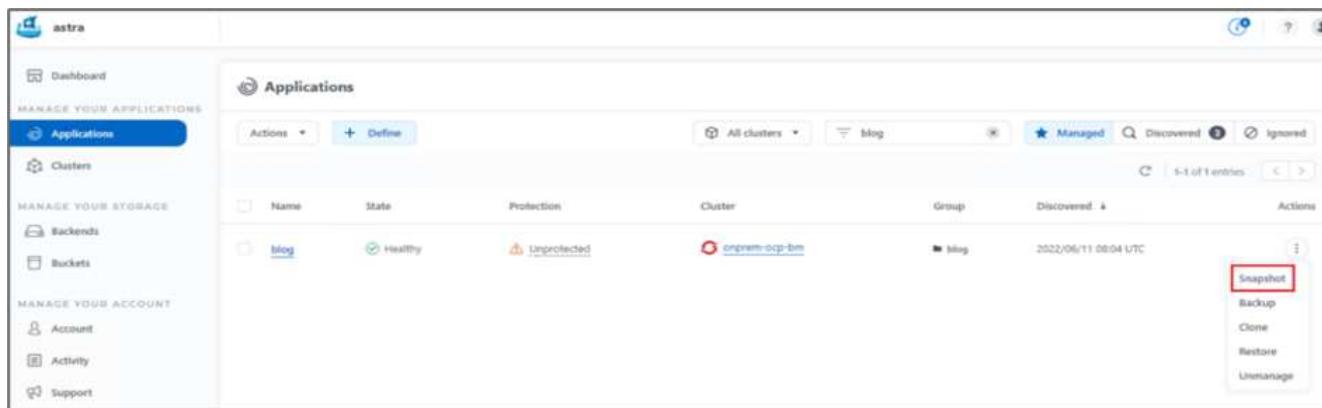


Astra Control

ABHINAV SINGH 11 JUN 2022 • 1 MIN READ

Astra Control is an application-aware data protection and mobility solution that manages, protects and moves data-rich Kubernetes workloads in both public clouds and on-premises. Astra Control enables data protection, disaster recovery, and migration for your Kubernetes workloads leveraging NetApp's industry-leading technology for snapshots, backups, replication, and cloning.

5. Go to Astra Control Center. First manage the app from the Discovered tab and then take a Snapshot copy.



The screenshot shows the Astra Control Center interface. On the left, there's a sidebar with 'Dashboard', 'Clusters', 'Rackends', 'Buckets', 'Account', 'Activity', and 'Support'. The main area is titled 'Applications' and shows a table with columns: Name, State, Protection, Cluster, Group, Discovered, and Actions. There is one entry: 'Blog' (State: Healthy, Protection: Unprotected, Cluster: onprem-ocp-bm, Group: blog, Discovered: 2022/06/11 08:04 UTC). The 'Actions' column for the 'Blog' entry has a red box around the 'Snapshot' button.



You can also protect your apps by creating snapshots, backups, or both at a defined schedule. For more information, see [Protect apps with snapshots and backups](#).

6. After the On-Demand snapshot is created successfully, upgrade the app to the latest version. The current image version is `ghost: 3.6-alpine` and the target version is `ghost:latest`. To upgrade the app, make changes directly to the Git repository and sync them to Argo CD.

```
spec:  
  containers:  
    - name: myblog  
      image: ghost:latest  
      imagePullPolicy: Always  
    ports:  
      - containerPort: 2368
```

7. You can see that the direct upgrade to the latest version is not supported due to the blog site being down and the entire application being corrupted.

Project: blog

Pods > Pod details

P myblog-5f899f7b76-zv7rq ● CrashLoopBackOff

Details Metrics YAML Environment Logs Events Terminal

Log stream ended. myblog Current log

```

34 lines
[2022-06-11 12:54:05] +[36mINFO+[39m Creating database backup
[2022-06-11 12:54:05] +[36mINFO+[39m Database backup written to: /var/lib/ghost/content/data/astra.ghost.2022-06-11-12-54-05.json
[2022-06-11 12:54:05] +[36mINFO+[39m Running migrations.
[2022-06-11 12:54:06] +[36mINFO+[39m Rolling back: Unable to run migrations.
[2022-06-11 12:54:06] +[36mINFO+[39m Rollback was successful.
[2022-06-11 12:54:06] +[31mERROR+[39m Unable to run migrations
+[31m
+[31mUnable to run migrations+[39m

+[37m"You must be on the latest v3.x to update across major versions - https://ghost.org/docs/update/"+[39m
+[33m"Run 'ghost update v3' to get the latest v3.x version, then run 'ghost update' to get to the latest."+[39m

+[1m+[37mError ID:+[39m+[22m
+[90m93b99ce0-e985-11ec-9301-7d29b2c73999+[39m

+[90m-----+[39m

+[90mInternalServerError: Unable to run migrations
    at /var/lib/ghost/versions/5.2.2/node_modules/knex-migrator/lib/index.js:1032:19
    at up (/var/lib/ghost/versions/5.2.2/core/server/data/migrations/utils/migrations.js:118:19)
    at Object.up (/var/lib/ghost/versions/5.2.2/core/server/data/migrations/utils/migrations.js:54:19)
    at /var/lib/ghost/versions/5.2.2/node_modules/knex-migrator/lib/index.js:982:33
    at /var/lib/ghost/versions/5.2.2/node_modules/knex/lib/execution/transaction.js:221:22+[39m
-[39m
[2022-06-11 12:54:06] +[35mWARN+[39m Ghost is shutting down
[2022-06-11 12:54:06] +[35mWARN+[39m Ghost has shut down
[2022-06-11 12:54:06] +[35mWARN+[39m Your site is now offline
[2022-06-11 12:54:06] +[35mWARN+[39m Ghost was running for a few seconds

```

8. To confirm the unavailability of the blog site, refresh the URL.

← → ⌂ Not secure | myblog-route-blog.apps.ocp.flexpod.netapp.com

Application is not available

The application is currently not serving requests at this endpoint. It may not have been started or is still starting.

i Possible reasons you are seeing this page:

- The host doesn't exist. Make sure the hostname was typed correctly and that a route matching this hostname exists.
- The host exists, but doesn't have a matching path. Check if the URL path was typed correctly and that the route was created using the desired path.
- Route and path matches, but all pods are down. Make sure that the resources exposed by this route (pods, services, deployment configs, etc) have at least one pod running.

9. Restore the app from the snapshot.

The screenshot shows the OpenShift Data Protection interface for the 'blog' application. The top navigation bar includes 'Actions' and a dropdown. The 'APPLICATION STATUS' section shows 'Healthy'. The 'APPLICATION PROTECTION STATUS' section shows 'Partially protected'. The 'Images' section lists 'ghost:3.6-alpine' and 'ghosttest'. The 'Protection schedule' is 'Disabled'. The 'Group' is 'Blog'. The 'Cluster' is 'onprem-ocp-bm'. The 'Data protection' tab is selected, showing a table with one entry: 'blog-snapshot-20220611125244'. The table columns are 'Name', 'State', 'On-Schedule / On-Demand', 'Created', and 'Actions'. The 'Actions' column for the snapshot shows 'Backup' (disabled), 'Restore application' (highlighted in red), and 'Delete snapshot'. A search bar and buttons for 'Snapshots' and 'Backups' are also present.

10. The app is restored on the same OpenShift cluster.

The screenshot shows the 'Restore namespace application' wizard, Step 2/2: SUMMARY. The title is 'Restore namespace application'. The sub-section is 'REVIEW RESTORE INFORMATION'. A warning message states: 'All existing resources associated with this namespace application will be deleted and replaced with the source snapshot "blog-snapshot-20220611125244" taken on 2022/06/11 12:52 UTC. Persistent volumes will be deleted and recreated. External resources with dependencies on this namespace application might be impacted.' It also says 'We recommend taking a snapshot or a backup of your namespace application before proceeding.' The left panel shows the 'SNAPSHOT' section with 'blog-snapshot-20220611125244'. The right panel shows the 'RESTORE' section with 'blog'. The 'DESTINATION GROUP' and 'DESTINATION CLUSTER' are both 'blog' and 'onprem-ocp-bm'. The 'RESOURCE LABELS' section shows 'Cluster Roles' and 'Cluster Role Bindings'. At the bottom, a confirmation message asks 'Are you sure you want to restore the namespace application "blog"?'. A text input field contains 'Confirm to restore' with 'restore' typed in. A 'Restore' button with a checkmark is at the bottom right.

11. The app restore process starts immediately.

The screenshot shows the 'Applications' list. The top navigation bar includes 'Actions', '+ Define', and search filters for 'All clusters', 'blog', 'Managed', 'Discovered', and 'Ignored'. The table shows one entry: 'blog' with a 'Restoring' state, 'Partially protected' protection, 'onprem-ocp-bm' cluster, 'blog' group, and 'Discovered' status. A search bar and a table header with columns 'Name', 'State', 'Protection', 'Cluster', 'Group', 'Discovered', and 'Actions' are also present.

12. In few minutes, the app is restored successfully from the available snapshot.

The screenshot shows the 'Applications' section of the Astra Control Center. At the top, there are buttons for 'Actions' (dropdown), '+ Define', and search fields for 'All clusters' (set to 'blog') and 'Managed' status. Below the search is a table header with columns: Name, State, Protection, Cluster, Group, Discovered, and Actions. A single row is listed: 'blog' is 'Healthy' (green checkmark), 'Partially protected' (blue info icon), and is associated with the 'onprem-ocp-bm' cluster, group 'blog', and was discovered on '2022/06/11 12:34 UTC'. There is a 'More' (three dots) button in the Actions column.

13. To see whether the webpage is available, refresh the URL.

The screenshot shows a web browser window with the URL 'Not secure | myblog-route-blog.apps.ocp.flexpod.netapp.com/astra-control-2/'. The page title is 'Astra Control'. Below the title, there is a user profile for 'ASHRAFAY DEV0H' and a timestamp '17 JUN 2022 - 1 MIN READ'. The main content area describes Astra Control as an application-aware data protection and mobility solution for Kubernetes workloads in both public clouds and on-premises. It highlights features like data protection, disaster recovery, and migration. At the bottom of the page is a navigation bar with icons for Home, Overview, Applications, and Support.

With the help of Astra Control Center, a DevTest team can successfully recover a blog site app and its associated data using the snapshot.

Part 2

With Astra Control Center, you can move an entire application along with its data from one Kubernetes cluster to another, no matter where the clusters are located (on-premises or in the cloud).

1. The DevTest team initially upgrades the app to the supported version (`ghost-4.6-alpine`) before upgrading to the final version (`ghost-latest`) to make it production ready. They then post an upgrade to the app that is cloned to the production OpenShift cluster running on a different FlexPod system.
2. At this point, the app is upgraded to the latest version and ready to be cloned to the production cluster.

Project: blog

Pods > Pod details

P myblog-55ffd9f658-tkbfq Running

Details Metrics **YAML** Environment Logs Events Terminal

```
180
181   ports:
182     - containerPort: 2368
183     protocol: TCP
184   imagePullPolicy: Always
185   volumeMounts:
186     - name: content
187       mountPath: /var/lib/ghost/content
188     - name: kube-api-access-t2sdz
189       readOnly: true
190       mountPath: /var/run/secrets/kubernetes.io/serviceaccount
191       terminationMessagePolicy: File
192       image: 'ghost:latest'
193   serviceAccount: default
194   volumes:
195     - name: content
196       persistentVolumeClaim:
197         claimName: blog-content
```

3. To verify the new theme, refresh the blog site.

Not secure | myblog-route-blog.apps.ocp.flexpod.netapp.com/astra-control-2/

Astra

Abhinav Singh
Jun 11, 2022

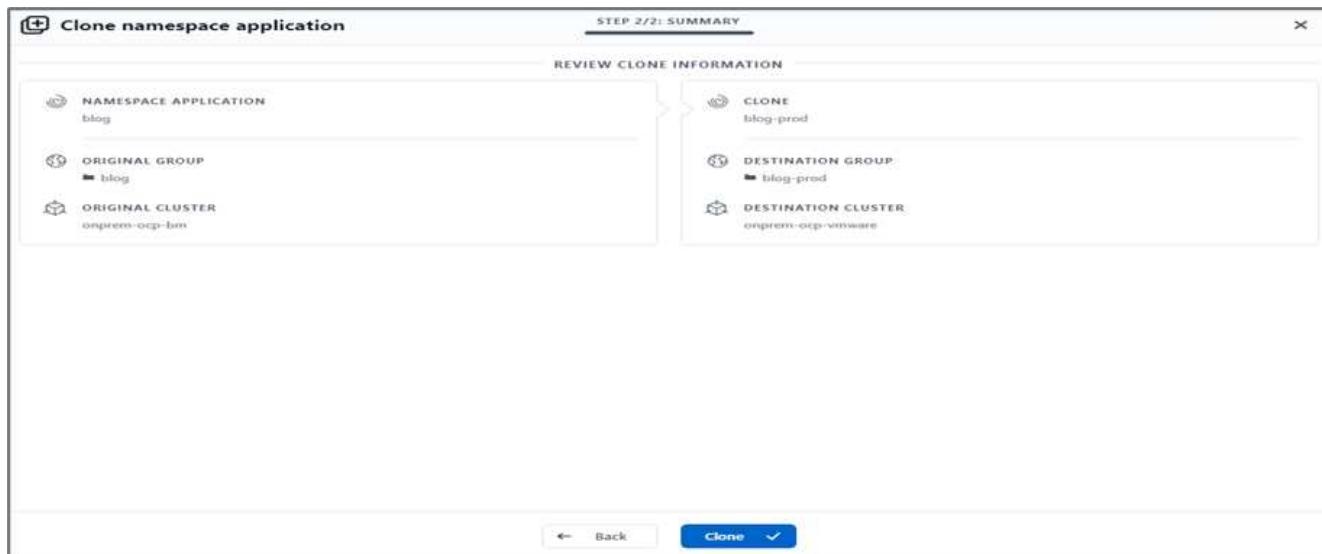
Astra Control

Astra Control is an application-aware data protection and mobility solution that manages, protects and moves data-rich Kubernetes workloads in both public clouds and on-premises. Astra Control enables data protection, disaster recovery, and migration for your Kubernetes workloads leveraging NetApp's industry-leading technology for snapshots, backups, replication, and cloning.

Sign up for more like this.

Enter your email **Subscribe**

4. From Astra Control Center, clone the app to the other production OpenShift cluster running on VMware vSphere.



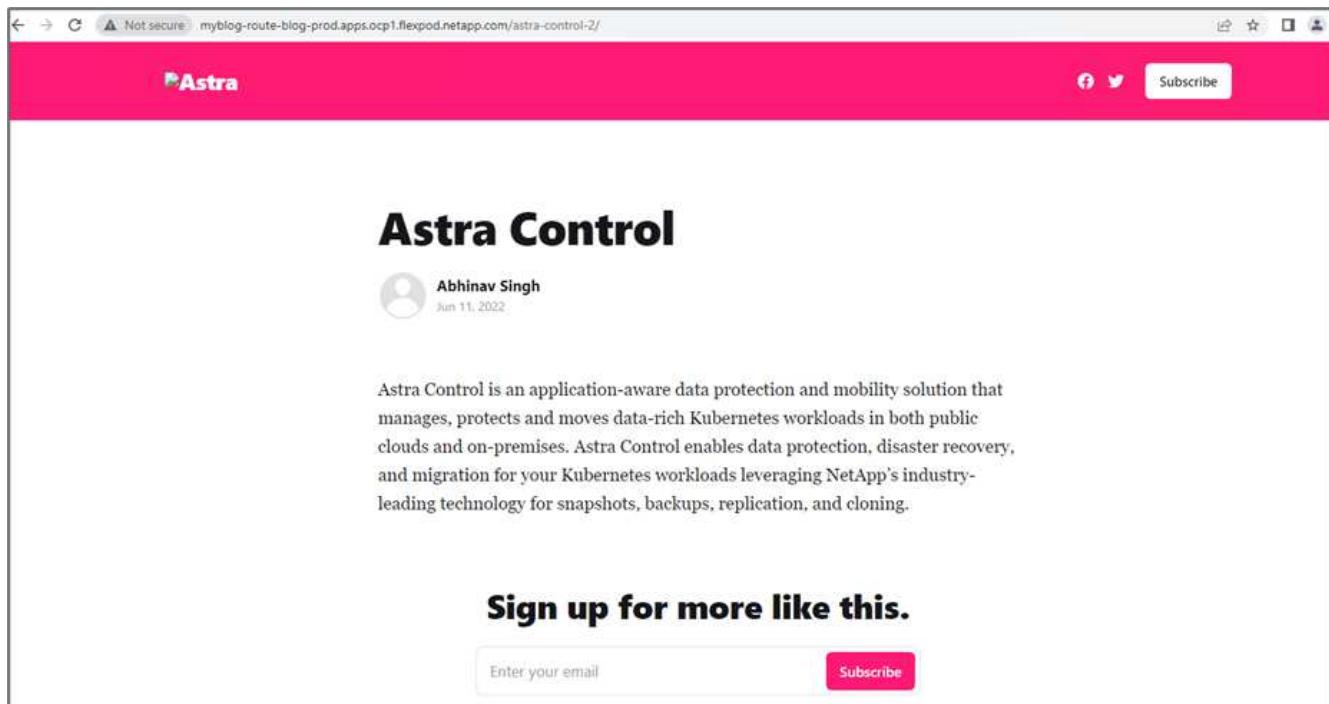
A new application clone is now provisioned in the production OpenShift cluster.

Name	State	Protection	Cluster	Group	Discovered	Actions
blog-prod	Provisioning	Unprotected	onprem-ocp-vmware	blog-prod	2022/06/11 13:17 UTC	⋮
blog	Healthy	Partially protected	onprem-ocp-bm	blog	2022/06/11 12:34 UTC	⋮

5. Log into the production OpenShift cluster and search for the project blog.

Name	Display name	Status	Requester	Memory	CPU	Created
PR blog-prod	No display name	Active	No requester	-	-	Just now

6. From the side menu, select Networking > Routes and click the URL under Location. The same homepage with the content is displayed.



The screenshot shows a blog post titled "Astra Control" by Abhinav Singh, published on Jun 11, 2022. The post discusses Astra Control as an application-aware data protection and mobility solution for Kubernetes workloads. It highlights features like data protection, disaster recovery, and migration. Below the post is a call-to-action section with a "Sign up for more like this." heading, an input field for an email address, and a "Subscribe" button.

Astra Control

 Abhinav Singh
Jun 11, 2022

Astra Control is an application-aware data protection and mobility solution that manages, protects and moves data-rich Kubernetes workloads in both public clouds and on-premises. Astra Control enables data protection, disaster recovery, and migration for your Kubernetes workloads leveraging NetApp's industry-leading technology for snapshots, backups, replication, and cloning.

Sign up for more like this.

Enter your email

This concludes the Astra Control Center solution validation. You can now clone an entire application and its data from one Kubernetes cluster to another no matter where the Kubernetes cluster is located.

[Next: Conclusion.](#)

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.