



NetApp Kubernetes Service: Accelerated Application Development and Deployment at Scale

EXECUTIVE SUMMARY

This guide shows how [NetApp® Kubernetes Service \(NKS\)](#) and [Cloud Volumes Service for AWS](#) can be used to deploy Kubernetes clusters with [persistent storage](#) in any of the major cloud environments. NetApp Build@Scale uses these services to create a software development platform that accelerates application development in hybrid environments. The paper describes the benefits of using NKS and the advantages of using Cloud Volumes Service for provisioning storage, and demonstrates how Build@Scale works. It also looks at the RESTful API interface for Cloud Volumes Service for AWS, and shows how it can be used to perform basic operations, such as creating volumes, snapshots, and clones.

INTRODUCTION

Kubernetes has become the de facto platform for container orchestration, dominating market share over competing solutions. All major cloud vendors now provide a managed Kubernetes service, such as [Amazon Elastic Kubernetes Service \(Amazon EKS\)](#) on AWS. However, customers can also choose to deploy Kubernetes with native cloud compute and storage resources, or by using on-premises systems. Orchestrating and managing Kubernetes deployments across such diverse environments can be extremely challenging, especially when persistent data storage is also a requirement.



NetApp Kubernetes Service (NKS) provides a single pane of glass from which you can deploy and manage Kubernetes clusters in AWS and hybrid cloud environments. NKS integrates with NetApp storage services, such as [Cloud Volumes Service](#), which allows the provisioning of Kubernetes persistent volumes to be powered by NetApp enterprise data storage solutions. [NetApp Kubernetes Service](#), therefore, gives users a complete solution for both infrastructure and application orchestration.

To accelerate the development of high-quality software applications, NetApp has also created Build@Scale, which runs on top of NKS and Cloud Volumes Service. This open-source solution gives software professionals a scalable and cohesive platform for managing DevOps pipelines and application developer workspaces.

This guide examines NetApp Kubernetes Service and Cloud Volumes Service in detail, and shows how Build@Scale uses these services to help you achieve faster time to market.

Amazon EKS

Amazon EKS is a cost-effective and fully managed solution for deploying Kubernetes clusters in AWS that removes the need for users to manage master nodes and the Kubernetes control plane. By deploying the Kubernetes management nodes across Availability Zones, clusters created with Amazon EKS are more resilient against service failures and outages.

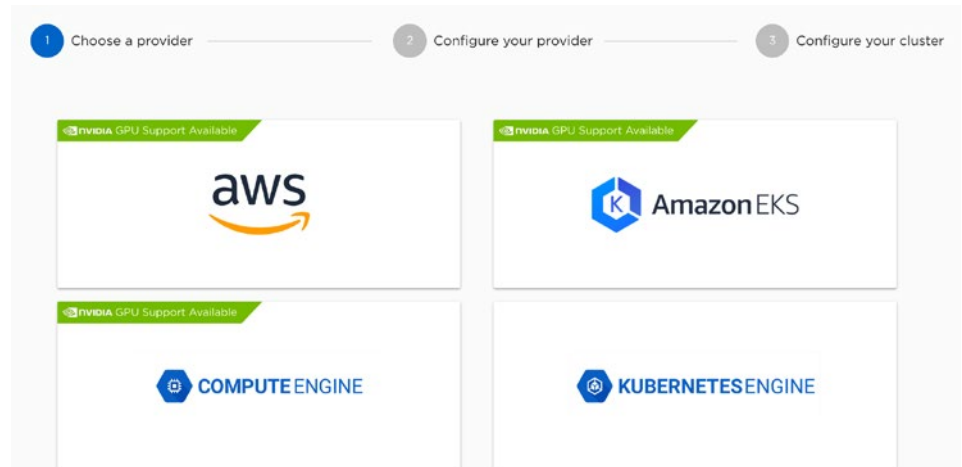
Worker nodes are deployed using [Amazon EC2](#), and all communication with the managed master nodes is encrypted by default. This makes it much easier to get up and running with Kubernetes in the cloud, because you can simply deploy the compute nodes you require without needing to stand up the base Kubernetes infrastructure and manage fault tolerance. Amazon EKS integrates easily with other AWS services, such as AWS CloudTrail for cluster logging and the mapping of Kubernetes RBAC roles directly to IAM entities.

Whether you're building out microservices or a batch processing platform, Amazon EKS gives you a reliable, flexible, and ready-to-use Kubernetes environment to start deploying your containers to. NetApp Kubernetes Service takes Amazon EKS to the next level by adding a suite of management tools and software packages that enhance your use of Kubernetes.

NetApp Kubernetes Service

NetApp Kubernetes Service provides uniformity when working with heterogeneous Kubernetes solutions, which may be on AWS. Support will soon extend to private cloud environments and

[NetApp HCI](#) as well. Whether you're using a managed Kubernetes service or building a cluster from the ground up, NKS gives you control over Kubernetes deployments and workloads in any scenario.

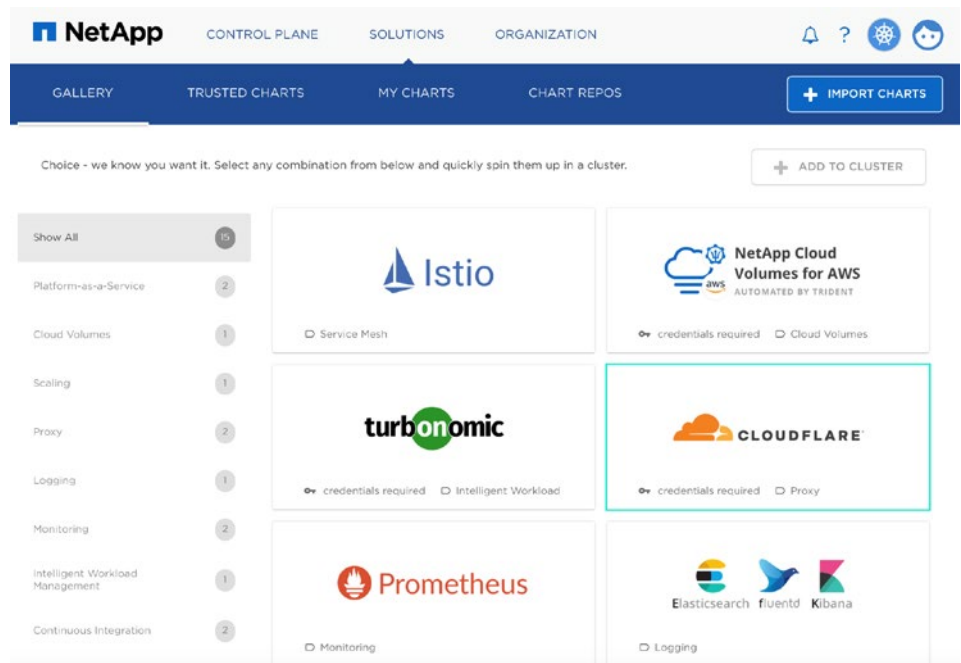


Deploying Kubernetes to AWS.

NKS is available as a software-as-a-service (SaaS) offering at [NetApp Cloud Central](#), so there's no software to install or configure. When creating a new cluster, a wizard-style interface guides you through the process of setting up the base infrastructure—such as the number and instance type of the master and worker nodes—and configuring Kubernetes. NKS then uses the information that has been provided to spin up a new, ready-to-use cluster within minutes.

Clusters deployed through NKS are very easy to keep up to date with the latest releases of Kubernetes. Users can perform an on-demand, zero-downtime upgrade to the latest version of Kubernetes, with a literal click of a button. This dramatically reduces the time administrators spend to deploy, maintain, and patch Kubernetes clusters.

The NetApp Kubernetes Service dashboard offers users a curated list of enterprise software solutions that can be installed directly into any of their clusters. These solutions help to manage, monitor, and scale the applications and services running in a cluster. For example, [Istio](#) simplifies the administration of a service mesh of microservices by providing facilities for managing cross-cluster service traffic rules, enforcing security policies and end-to-end data encryption, logging and service observability, and much more.



Easily deploy trusted applications into your cluster.

NetApp Kubernetes Service offers many other benefits, including:

- **Service management.** NKS makes it easier to deploy and manage microservices in multiregion and hybrid environments by providing first-rate support for solutions such as Istio.
- **Node pooling.** Users can create individual pools of nodes with different resource profiles. Each pool can also be independently autoscaled.
- **Master scaling.** Master nodes can be easily scaled in both size and number. Having more than one master node avoids a single point of failure for Kubernetes clusters.
- **Application upgrades.** NKS supports canary and blue/green deployments for releasing new versions of applications and services without downtime.
- **Role-based access control.** Administrators can control access to NKS by using a set of predefined user roles.
- **Backup and recovery.** NKS supports the backup and recovery of all Kubernetes configuration information, as well as persistent volume data.

Cloud Volumes Service and NetApp Trident

NetApp Kubernetes Service offers direct support for dynamically provisioning cluster storage by using NetApp storage solutions. This support is accomplished through [NetApp Trident](#), an open-source Kubernetes storage provisioner that provides integration with any ONTAP® storage environment, including [Cloud Volumes ONTAP](#) and [Cloud Volumes Service for AWS](#).

SERVICE LEVELS

Users can change their performance level on the fly to get the right performance at the right time.

Standard service level. Up to 1000 IOPS per TB (16k I/O) and 16 MB/s of throughput per TB.

Premium service level. Up to 4000 IOPS per TB (16k I/O) and 64 MB/s of throughput per TB.

Extreme service level. Up to 8000 IOPS per TB (16k I/O) and 128 MB/s of throughput per TB.

Cloud Volumes Service is a fully managed file share platform with which users can deploy enterprise-grade NFS and SMB storage volumes with a single click. Every deployed file share benefits from:

- **High performance.** Configure file share volumes with a service level that determines the I/O performance they deliver, with a choice of Standard, Premium, or Extreme.
- **Scalability.** Nondisruptively scale up performance in response to increased concurrent load.
- **High availability.** Cloud Volumes Service is built with redundancy at every level in the stack, which ensures that file shares are available and all data is durable.
- **Data security.** All data is encrypted at rest for transparent protection against unauthorized access.
- **Snapshots.** Instantly create point-in-time copies of a file share volume for data protection and versioning. Snapshots can also be instantly restored back to the source volume.
- **Cloning.** Use snapshots to rapidly provision new file share volumes based on the data in existing volumes. This is highly advantageous for software development testing and setting up DevOps pipelines.
- **Data synchronization.** Copy data into file shares from external systems, or vice versa, while keeping data incrementally synchronized going forward.

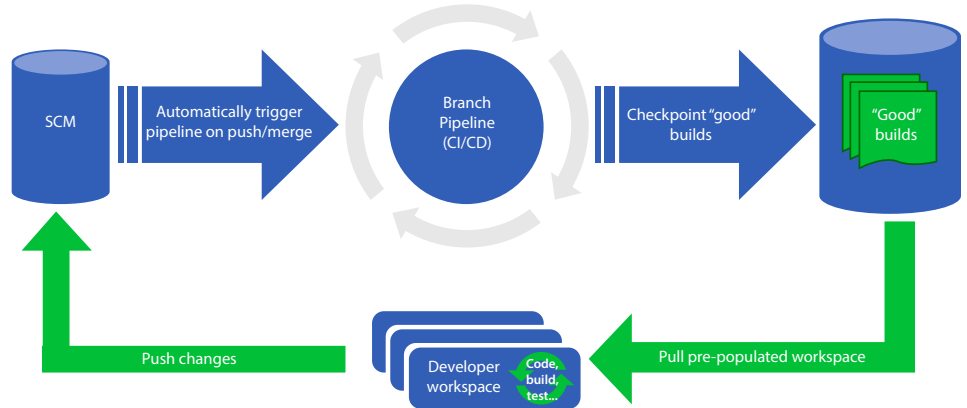
NetApp Kubernetes Service automatically deploys NetApp Trident into new cluster nodes, which allows storage to be provisioned for NFS persistent volume claims, in whichever cloud you are using. This removes the complexity of provisioning persistent storage for Kubernetes clusters in hybrid cloud environments and massively reduces administrative overhead. The feature sets of NKS and Cloud Volumes Service can also be used to drive application development with Build@Scale, discussed in the next section.

Application Development with Build@Scale

[Build@Scale](#) is another open-source solution from NetApp that can use NetApp Kubernetes Service and Cloud Volumes Service to give software developers a platform for creating high-quality applications in the cloud. This is achieved by incorporating [DevOps](#) pipelines and developer workspaces into a single Kubernetes cluster, with Cloud Volumes Service providing persistent storage, snapshots, and cloning facilities.

Build@Scale preinstalls pods for various applications, including Gitlab for source control and Jenkins for managing DevOps jobs and processes. To create a new DevOps pipeline, simply give Build@Scale the URL of a source code repository and the branch to be used. Build@Scale


then creates and configures a new Jenkins job for the pipeline, with a corresponding file share volume for supporting incremental builds.



The Build@Scale pipeline overview.

Whenever the pipeline is executed, the latest code is pulled from source control and then built and compiled. A storage snapshot is made of the resulting pipeline workspace, which is then used to clone developer workspaces. Developers can access these workspaces through a web-based integrated development environment, such as [Theia](#).

Stage View



Execution of a DevOps CI pipeline.

Using Build@Scale, you can have many pipelines active at the same time, either for different applications and source code repositories or for different branches of the same repository, with each pipeline supporting multiple concurrent developer workspaces. This allows you to scale the platform in tandem with your pace of application development, while at the same time improving software quality by automatically triggering continuous integration (CI) builds when source code is pushed.

Cloud Volumes Service REST API

Build@Scale uses the RESTful API provided by Cloud Volumes Service to create new file share volumes, to make snapshots those volumes, and to create clones from the snapshots. NetApp Trident uses the same API to provision persistent volumes in a Kubernetes cluster. End users

API INTERFACE

Makes it possible to:

- Create new volumes
- Manage snapshots
- Create clones
- And much more

who require greater flexibility when working with file share volumes can also access the Clouds Volumes Service REST API directly.

Accessing the API

To access the Cloud Volumes Service API, you must first find the URL for the API in your deployment environment, as well as the API key and Secret key that will be used for authentication. All of this information is available from the Cloud Volumes Service main dashboard.

Calling API Endpoints

The Cloud Volumes Service dashboard also contains a link for the API documentation, which provides detailed information about each available endpoint, including the parameters and message payload it expects and the data it returns. Users can even try out the endpoints from the API documentation portal.

The following are simple examples of calling the API using cURL. In each case, the API URL, API key, and Secret key placeholders should be replaced with the values for your environment.

Creating a new volume

```
curl -H "accept:application/json" \  
      -H "Content-type: application/json" \  
      -H "api-key:<api_key>" \  
      -H "secret-key:<secret_key>" \  
      -X POST <api_url>/v1/FileSystems -d '{  
  "name": "Test",  
  "creationToken": "test-volume1",  
  "region": "us-west",  
  "serviceLevel": "basic",  
  "quotaInBytes": 100000000000,  
  "exportPolicy": {  
    "rules": [  
      {  
        "ruleIndex": 1,  
        "allowedClients": "0.0.0.0/0",  
        "unixReadOnly": false,  
        "unixReadWrite": true,  
        "cifs": false,  
      }  
    ]  
  }  
'
```

```

        "nfsv3": true,
        "nfsv4": false
    }
]
},
"labels": [
    "test"
]
}'

```

Managing Snapshots

For snapshots, the placeholder for the ID of the volume to operate on must also be specified. The following command can be used to create a new snapshot:

```

curl -H "accept:application/json" \
-H "Content-type: application/json" \
-H "api-key:<api_key>" \
-H "secret-key:<secret_key>" \
-X POST <api_url>/v1/FileSystems/<volume_
id>/Snapshots -d '{
    "name": "snap",
    "region": "us-west"
}'

```

The following command can be used to list all snapshots for a volume, which can be useful when creating volume clones:

```

curl -H "accept:application/json" \
-H "Content-type: application/json" \
-H "api-key:<api_key>" \
-H "secret-key:<secret_key>" \
-X GET <api_url>/v1/
FileSystems/<volume_id>/Snapshots

```


Creating a Clone

Clones are created by using snapshots of existing file share volumes. The ID of the snapshot to use should be provided in the following command:

```
curl -H "accept:application/json" \
      -H "Content-type: application/json" \
      -H "api-key:<api_key>" \
      -H "secret-key:<secret_key>" \
      -X POST <api_url>/v1/FileSystems -d '{
"snapshotId": "<snapshot_id>",
"name": "testClone",
"creationToken": "my-test-clone",
"region": "us-west",
"serviceLevel": "extreme"
}'
```

Conclusion

NetApp Kubernetes Service and Cloud Volumes Service for AWS provide an end-to-end solution for deploying Kubernetes with persistent volumes in AWS and hybrid cloud environments. This empowers customers to fully unleash the scalability of all cloud environments without being impeded by the administration of cloud compute and storage infrastructure.

Build@Scale uses NKS and Cloud Volumes Service to create a scalable application development platform in any cloud. By simplifying the setup of DevOps pipelines and developer workspaces, Build@Scale helps to accelerate application development by improving developer velocity. When your application is ready for production, NKS takes care of the rollout.

Visit [NetApp](#) online to find out more about speeding up your application development lifecycle by using [NetApp Kubernetes Service](#), [Cloud Volumes Service for AWS](#), and [Build@Scale](#). Or get started now: [NetApp Kubernetes Service](#) and [NetApp Cloud Volumes Service for AWS](#) are available on the [AWS Marketplace](#).