



FlexPod Hybrid Cloud for Google Cloud Platform with NetApp Cloud Volumes ONTAP and Cisco Intersight

FlexPod

NetApp
January 21, 2025

This PDF was generated from <https://docs.netapp.com/us-en/flexpod/hybrid-cloud/gcp-ncvo-solution-overview.html> on January 21, 2025. Always check docs.netapp.com for the latest.

Table of Contents

FlexPod Hybrid Cloud for Google Cloud Platform with NetApp Cloud Volumes ONTAP and Cisco Intersight . . 1

TR-4939: FlexPod Hybrid Cloud for Google Cloud Platform with NetApp Cloud Volumes ONTAP and Cisco Intersight. 1

Solution components 3

Installation and configuration 8

Solution validation 72

Conclusion 80

FlexPod Hybrid Cloud for Google Cloud Platform with NetApp Cloud Volumes ONTAP and Cisco Intersight

TR-4939: FlexPod Hybrid Cloud for Google Cloud Platform with NetApp Cloud Volumes ONTAP and Cisco Intersight

Ruchika Lahoti, NetApp

Introduction

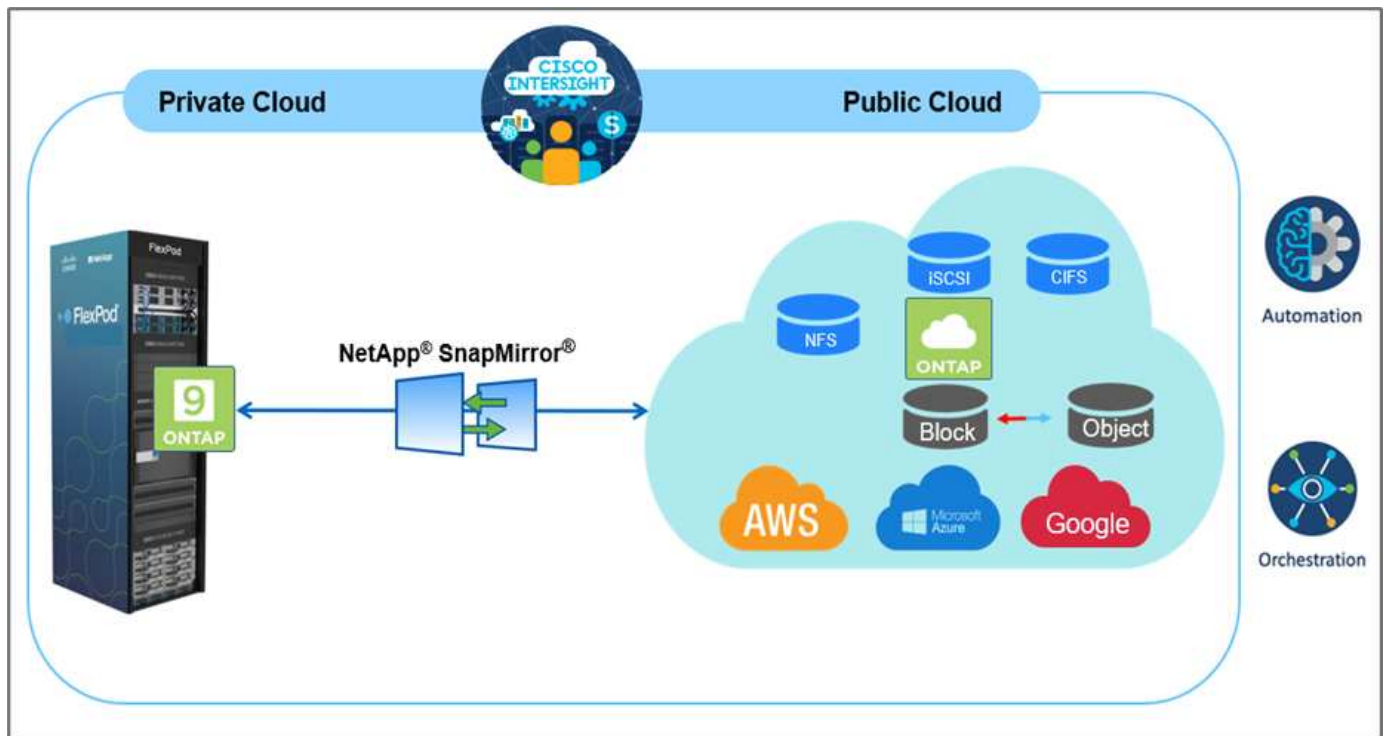
Protecting data with disaster recovery (DR) is a critical goal for businesses continuity. DR allows organizations to failover their business operations to a secondary location and later recover and failback to the primary site efficiently and reliably. Multiple concerns like natural disaster, network failures, software vulnerabilities, and human error make developing a DR strategy a top IT priority.

For DR, all workloads running on the primary site must be faithfully reproduced on the DR site. An organization must also have an up-to-date copy of all enterprise data, including database, file services, NFS and iSCSI storage, and so on. Because data in the production environment is constantly updated, changes must be transferred to the DR site on a regular basis.

Deploying DR environments is challenging for most organizations because of the requirement for infrastructure and site independence. The number of resources needed and the costs of setting up, testing, and maintaining a secondary data center can be very high, typically approaching the cost of the entire production environment. It is challenging to keep a minimal data footprint with adequate protection, while continuously synchronizing data and establishing seamless failover and failback. After building out the DR site, the challenge then becomes to replicate data from the production environment and to keep it synchronized going forward.

This technical report brings together the FlexPod converged infrastructure solution, NetApp Cloud Volumes ONTAP on Google Cloud, and Cisco Intersight to form a hybrid cloud data center for DR. In this solution we discuss designing and executing an on-premises ONTAP workflow using Cisco Intersight Cloud Orchestrator. We also discuss deploying NetApp Cloud Volumes ONTAP and orchestrating and automating data replication and DR between FlexPod and Cloud Volumes ONTAP using the Cisco Intersight Service for HashiCorp Terraform.

The following figure provide a solution overview.



This solution provides multiple advantages, including:


- **Orchestration and automation.** Cisco Intersight simplifies the day-to-day operations of FlexPod hybrid cloud infrastructure by providing consistent orchestration frameworks that are delivered via automation.
- **Customized Protection.** Cloud Volumes ONTAP provides block-level data replication from ONTAP to the cloud that keeps the destination up to date through incremental updates. Users can specify a synchronization schedule of every 5 minutes or every hour, for example, based on changes at the source that are transferred over.
- **Seamless failover and failback.** When a disaster occurs, storage administrators can quickly failover to the cloud volumes. When the primary site is recovered, the new data created in the DR environment is synchronized back to the source volumes, re-establishing secondary data replication.
- **Efficiency:** The storage space and costs for the secondary cloud copy are optimized through the use of data compression, thin provisioning, and deduplication. Data is transferred at the block-level in a compressed and deduplicated form, improving transfer speed. Data is also automatically tiered to low-cost object storage and only brought back to high-performance storage when accessed, such as in a DR scenario. This significantly reduces ongoing storage costs.
- **Increased IT productivity.** Using Intersight as the single secure, enterprise-grade platform for infrastructure and application lifecycle management simplifies configuration management and automation of manual tasks at scale for the solution.

Audience

The audience for this document includes, but is not limited to; sales engineers, field consultants, professional services, IT managers, partner engineers, site reliability engineers, cloud architects, cloud engineers, and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

Solution topology

This section describes the logical topology of the solution. The following figure represents the solution topology of the on-premises FlexPod environment, NetApp Cloud Volumes ONTAP running on Google Cloud, Cisco Intersight, and NetApp Cloud Manager.



The control planes and data planes are clearly indicated between the endpoints. The data plane uses a secure site-to-site VPN connection to connect the ONTAP instance running on FlexPod All Flash FAS to the NetApp Cloud Volumes ONTAP instance on Google Cloud.

The replication of workload data from FlexPod to NetApp Cloud Volumes ONTAP is handled by NetApp SnapMirror, and the overall process is orchestrated using Cisco Intersight Cloud Orchestrator for both the on-premises and cloud environments. Cisco Intersight Cloud Orchestrator consumes Terraform Resource Providers for NetApp Cloud Manager to carry out operations related to NetApp Cloud Volumes ONTAP deployment and establish data replication relationships.



The optional backup and tiering of cold data residing in the NetApp Cloud Volumes ONTAP instance to Google Cloud Storage is also supported with this solution.

[Next: Solution components.](#)

Solution components

[Previous: Solution overview.](#)

FlexPod

FlexPod is a defined set of hardware and software that forms an integrated foundation for both virtualized and non-virtualized solutions. FlexPod includes NetApp ONTAP storage, Cisco Nexus networking, Cisco MDS storage networking, and Cisco Unified Computing System (Cisco UCS). The design is flexible enough that the networking, computing, and storage can fit into one data center rack, or it can be deployed according to a

customer's data center design. Port density allows the networking components to accommodate multiple configurations.

Cisco Intersight

Cisco Intersight is a SaaS platform that delivers intelligent automation, observability, and optimization for traditional and cloud-native applications and infrastructure. The platform helps to drive change with IT teams and delivers an operating model designed for hybrid cloud. Cisco Intersight provides the following benefits:

- **Faster delivery.** Delivered as a service from the cloud or in the customer's data center with frequent updates and continued innovation, due to an agile-based software development model. This way the customer can focus on accelerating delivery for line-of-business.
- **Simplified operations.** Simplify operations by using a single secure SaaS-delivered tool with common inventory, authentication, and APIs to work across the full stack and all locations, eliminating silos across teams. From managing physical servers and hypervisors on-premises, to VMs, K8s, serverless, automation, optimization, and cost control across both on-premises and public clouds.
- **Continuous optimization.** Continuously optimize your environment by using intelligence provided by Cisco Intersight across every layer, as well as Cisco TAC. This intelligence is converted into recommended and automatable actions so you can adapt real-time to every change: from moving workloads and monitoring health of physical servers to cost reduction recommendations the public clouds you work with.

There are two modes of management operations possible with Cisco Intersight: UCSM Managed Mode (UMM) and Intersight Managed Mode (IMM). You can select native UMM or IMM for fabric-attached Cisco UCS systems during initial setup of fabric interconnects. In this solution, native IMM is used.

Cisco Intersight licensing

Cisco Intersight uses a subscription-based license with multiple tiers.

Cisco Intersight license tiers are as follows:

- **Cisco Intersight Essentials.** Includes all base functionality plus the following features:
 - Cisco UCS Central
 - Cisco IMC Supervisor entitlement
 - Policy-based configuration with Server Profiles
 - Firmware management
 - Valuation of compatibility with the Hardware Compatibility List (HCL)
- **Cisco Intersight Advantage.** Includes of the features and functionality of the Essentials tier plus the following features:
 - Widgets, inventory, capacity, utilization features, and cross-domain inventory correlation across physical compute, network, storage, VMware virtualization, and AWS public cloud.
 - The Cisco Security Advisory service where customers can receive important security alerts and field notices about impacted endpoint devices.
- **Cisco Intersight Premier.** In addition to the capabilities provided in the Advantage tier, Cisco Intersight Premier offers the following:
 - Intersight Cloud Orchestrator (ICO) for Cisco and third-party compute, network, storage, integrated systems, virtualization, container, and public-cloud platforms
 - Full subscription entitlement for Cisco UCS Director at no additional cost.

More information about Intersight Licensing and features supported in each licensing can be found [here](#).



In this solution, we use Intersight Cloud Orchestrator and Intersight Service for HashiCorp Terraform. These features are available for users with the Intersight Premier license, so this licensing tier must be enabled.

Terraform Cloud Integration with ICO

You can use Cisco Intersight Cloud Orchestrator (ICO) to create and execute workflows that call Terraform Cloud (TFC) APIs. The Invoke Web API Request task supports Terraform Cloud as a target, and it can be configured with Terraform Cloud APIs using HTTP methods. So, the workflow can have a combination of tasks that calls multiple Terraform Cloud APIs using generic API tasks and other operations. You need a Premier license to use the ICO feature.

Cisco Intersight Assist

Cisco Intersight Assist helps you add endpoint devices to Cisco Intersight. A data center could have multiple devices that do not connect directly with Cisco Intersight. Any device that is supported by Cisco Intersight but does not connect directly to it requires a connection mechanism. Cisco Intersight Assist provides that connection mechanism and helps you add devices into Cisco Intersight.

Cisco Intersight Assist is available within the Cisco Intersight Virtual Appliance, which is distributed as a deployable virtual machine contained within an Open Virtual Appliance (OVA) file format. You can install the appliance on an ESXi server. For more information, see the [Cisco Intersight Virtual Appliance Getting Started Guide](#).

After claiming Intersight Assist into Intersight, you can claim endpoint devices using the Claim Through Intersight Assist option. For more information, see [Getting Started](#).

NetApp Cloud Volumes ONTAP

- Leveraging built-in data deduplication, data compression, thin provisioning, and cloning to minimize storage costs.
- Providing enterprise reliability and continuous operations in case of failures in your cloud environment.
- Cloud Volumes ONTAP uses NetApp SnapMirror, industry-leading replication technology, to replicate on-premises data to the cloud so it's easy to have secondary copies available for multiple use cases.
- Cloud Volumes ONTAP also integrates with the Cloud Backup service to deliver backup and restore capabilities for protection and long-term archiving of your cloud data.
- Switching between high and low-performance storage pools on-demand without taking applications offline.
- Providing consistency of Snapshot copies using NetApp SnapCenter.
- Cloud Volumes ONTAP supports data encryption and provides protection against viruses and ransomware.
- Integration with Cloud Data Sense helps you understand data context and identify sensitive data.

Cloud Central

Cloud Central provides a centralized location to access and manage NetApp cloud data services. These services enable you to run critical applications in the cloud, create automated DR sites, back up your SaaS data, and effectively migrate and control data across multiple clouds. For more information, see [Cloud Central](#).

Cloud Manager

Cloud Manager is an enterprise-class, SaaS-based management platform that enables IT experts and cloud architects to centrally manage their hybrid multi-cloud infrastructure using NetApp cloud solutions. It provides a centralized system for viewing and managing your on-premises and cloud storage to support multiple hybrid-cloud providers and accounts. For more information, see [Cloud Manager](#).

Connector

Connector enables Cloud Manager to manage resources and processes within a public cloud environment. A Connector instance is required to use many features provided by Cloud Manager and can be deployed in the cloud or on-premises network. Connector is supported in the following locations:

- AWS
- Microsoft Azure
- Google Cloud
- On premises

NetApp Active IQ Unified Manager

NetApp Active IQ Unified Manager allows you to monitor your ONTAP storage clusters from a single, redesigned, intuitive interface that delivers intelligence from community wisdom and AI analytics. It provides comprehensive operational, performance, and proactive insights into the storage environment and the virtual machines running on it. When an issue occurs with the storage infrastructure, Unified Manager can notify you about the details of the issue to help identify the root cause. The virtual machine dashboard gives you a view into the performance statistics for the VM so that you can investigate the entire I/O path from the vSphere host down through the network and finally to the storage.

Some events also provide remedial actions that you can take to rectify the issue. You can configure custom alerts for events so that when issues occur, you are notified through email and SNMP traps. Active IQ Unified Manager enables planning for the storage requirements of your users by forecasting capacity and usage trends to proactively act before issues arise, preventing reactive short-term decisions that can lead to additional problems in the long term.

VMware vSphere

VMware vSphere is a virtualization platform for holistically managing large collections of infrastructures (resources including CPUs, storage, and networking) as a seamless, versatile, and dynamic operating environment. Unlike traditional operating systems that manage an individual machine, VMware vSphere aggregates the infrastructure of an entire data center to create a single powerhouse with resources that can be allocated quickly and dynamically to any application in need.

For more information about VMware vSphere, follow [this link](#).

VMware vSphere vCenter

VMware vCenter Server provides unified management of all hosts and VMs from a single console and aggregates performance monitoring of clusters, hosts, and VMs. VMware vCenter Server gives administrators a deep insight into the status and configuration of compute clusters, hosts, VMs, storage, the guest OS, and other critical components of a virtual infrastructure. VMware vCenter manages the rich set of features available in a VMware vSphere environment.

Hardware and software versions

This hybrid cloud solution can be extended to any FlexPod environment that is running supported versions of software, firmware, and hardware as defined in the NetApp Interoperability Matrix Tool and the Cisco UCS Hardware Compatibility List.

The FlexPod solution that is used as a baseline platform in our on-premises environment was deployed according to the guidelines and specifications described [here](#).

The network within this environment is ACI- based. For more information, see [here](#).

- See the following links for more information:
- [NetApp Interoperability Matrix Tool](#)
- [VMware Compatibility Guide](#)
- [Cisco UCS Hardware and Software Interoperability Tool](#)

The following table shows the FlexPod hardware and software revisions.

Component	Product	Version
Compute	Cisco UCS X210C-M6	5.0(1b)
	Cisco UCS Fabric Interconnects 6454	4.2(2a)
Network	Cisco Nexus 9332C (Spine)	14.2(7s)
	Cisco Nexus 9336C-FX2 (Leaf)	14.2(7s)
	Cisco ACI	4.2(7s)
Storage	NetApp AFF A220	9.11.1
	NetApp ONTAP Tools for VMware vSphere	9.10
	NetApp NFS Plugin for VMware VAAI	2.0-15
	Active IQ Unified Manager	9.11
Software	vSphere ESXi	7.0(U3)
	VMware vCenter Appliance	7.0.3
	Cisco Intersight Assist Virtual Appliance	1.0.11-306

The execution of Terraform configurations happens on the Terraform Cloud for Business account. Terraform configuration uses the Terraform provider for NetApp Cloud Manager.

The following table lists the vendors, products, and versions.

Component	Product	Version
HashiCorp	Terraform	1.2.7

The following table shows the Cloud Manager and Cloud Volumes ONTAP versions.

Component	Product	Version
NetApp	Cloud Volumes ONTAP	9.11
	Cloud Manager	3.9.21

[Next: Installation and configuration - Deploy FlexPod.](#)

Installation and configuration

Deploy FlexPod

[Previous: Solution components.](#)

To understand the FlexPod design and deployment details, including the configuration of various elements of design and associated best practices, see [Cisco Validated Designs for FlexPod](#).

FlexPod can be deployed in both UCS Managed Mode and Cisco Intersight Managed Mode. If you are deploying FlexPod in UCS Managed Mode, the latest Cisco Validated Design can be found [here](#).

Cisco Unified Compute System (Cisco UCS) X-Series is a brand new modular compute system, configured and managed from the cloud. It is designed to meet the needs of modern applications and to improve operational efficiency, agility, and scale through an adaptable, future-ready, modular design. The design guidance around incorporating the Cisco Intersight- managed UCS X-Series platform within FlexPod infrastructure can be found [here](#).

FlexPod with Cisco ACI deployment can be found [here](#).

[Next: Cisco Intersight configuration.](#)

Cisco Intersight configuration

[Previous: Deploy FlexPod.](#)

To configure Cisco Intersight and Intersight Assist, see the Cisco Validated Designs for FlexPod found [here](#).

[Next: Terraform Cloud Integration with ICO prerequisite.](#)

Terraform Cloud Integration with ICO prerequisite

[Previous: Cisco Intersight configuration.](#)

Procedure 1: Connect Cisco Intersight and Terraform Cloud

1. Claim or create a Terraform cloud target by providing the relevant Terraform Cloud account details.
2. Create a Terraform Cloud Agent target for private clouds so that customers can install the agent in the data center and enable communication with Terraform Cloud.

For more information, follow [this link](#).


Procedure 2: Generate user token

As a part of adding a target for Terraform Cloud, you must provide the username and API token from the Terraform Cloud settings page.

1. Login to Terraform Cloud and go to **User Tokens**: <https://app.terraform.io/app/settings/tokens>.
2. Click **Create a new API token**.
3. Assign a name to remember and save the token in a secure place.

Procedure 3: Claim Terraform Cloud Target

1. Log into Intersight with Account Administrator, Device Administrator, or Device Technician privileges.
2. Navigate to **ADMIN > Targets > Claim a New Target**.
3. In **Categories**, click **Cloud**.
4. Click **Terraform Cloud** and click **Start**.



5. Enter a name for the target, your username for the Terraform Cloud, the API token, and a default organization in Terraform Cloud as displayed in the following image.
6. In the **Default Managed Hosts** field, make sure to add the following links along with other managed hosts:
 - github.com
 - github-releases.githubusercontent.com

The screenshot shows a configuration form for a Terraform Cloud target. It has a dark blue background with white text. The form contains four fields arranged in two rows. The first row has 'Name *' with the value 'TFCB' and 'Terraform Cloud API Token' with a masked value '*****'. The second row has 'Terraform Cloud Username *' with the value 'abhinav3' and 'Default Managed Hosts' with the value 'github.com,github-releases.githubusercontent.com'. Each field has a small circular icon to its right. The 'Terraform Cloud API Token' field also has a pencil icon for editing.

Name *	Terraform Cloud API Token
TFCB	*****
Terraform Cloud Username *	
abhinav3	
Default Terraform Cloud Organization *	Default Managed Hosts
cisco-intersight-gc	github.com,github-releases.githubusercontent.com

If everything is correctly entered, you will see your Terraform Cloud target displayed in the **Intersight Targets** section.

Procedure 4: Add Terraform Cloud agents

Prerequisites:

- Terraform Cloud target.
- Claimed Intersight Assist into Intersight before deploying the Terraform Cloud Agent.



You can only claim five agents for each Assist.



After you have created the connection to Terraform, you must spin up a Terraform Agent to execute the Terraform code.

1. Click **Claim Terraform Cloud Agent** from the drop-down list of your Terraform Cloud target.
2. Enter the details for the Terraform Cloud agent. The following screenshot shows the configuration details for Terraform agent.

Terraform Cloud target

Name *
flexpod-solution-terraform-agent

Intersight Assist *
g13-intersight-appliance.fpmc.sa

Terraform Cloud Organization *
cisco-intersight-gc

Terraform Cloud Agent Pool Name *
flexpod-solution-agent-pool

Managed Hosts

Hostname / IP Address / Subnets *	
github.com	🗑️
github-releases.githubusercontent.com	🗑️

+



You can update any Terraform Agent property. If the target is in the **Not Connected** state and has never been in the **Connected** state, then a token has not been generated for the Terraform agent.

After the agent validation succeeds and an agent token is generated, you are unable to reconfigure the Organization and/or Agent Pool. Successful deployment of a Terraform agent is indicated by a status of **Connected**.

After you have enabled and claimed the Terraform Cloud integration, you can deploy one or more Terraform Cloud agents in Cisco Intersight Assist. The Terraform Cloud agent is modelled as a child target of the Terraform Cloud target. When you claim the agent target, you see a message to indicate that the target claim is in progress.

After a few seconds, the target is moved to the **Connected** state, and the Intersight platform routes HTTPS packets from the agent to the Terraform Cloud gateway.

Your Terraform Agent should be correctly claimed and should show up under targets as **Connected**.

[Next: Configure Public Cloud Service provider.](#)

Configure Public Cloud Service provider

[Previous: Terraform Cloud Integration with ICO prerequisite.](#)

Procedure 1: Access NetApp Cloud Manager

To access NetApp Cloud Manager and other cloud services, you need to sign up on [NetApp Cloud Central](#).



For setting up workspaces and users in the Cloud Central account, click [here](#).

Procedure 2: Deploy Connector

To deploy Connector in Google Cloud, see this [link](#).

[Next: Automated deployment of Hybrid Cloud NetApp Storage.](#)

Automated deployment of Hybrid Cloud NetApp Storage

[Previous: Configure Public Cloud Service provider.](#)

Google Cloud

You must first enable APIs and create a service account that provides Cloud Manager with permissions to deploy and manage Cloud Volumes ONTAP systems that are in the same project as the Connector or in different projects.

Before you deploy a connector in a Google Cloud project, make sure that the connector isn't running on your premises or in a different cloud provider.

Two sets of permissions must be in place before you deploy a Connector directly from Cloud Manager:

- You need to deploy Connector using a Google account that has permissions to launch the Connector VM instance from Cloud Manager.
- When deploying Connector, you are prompted to select the VM instance. Cloud Manager gets permissions from the service account to create and manage Cloud Volumes ONTAP systems on your behalf. Permissions are provided by attaching a custom role to the service account. You need to set up two YAML files that include the required permissions for the user and the service account. Learn how to use [the YAML files to set up permissions](#) here.

See [this detailed video](#) for all required prerequisites.

Cloud Volumes ONTAP deployment modes and architecture

Cloud Volumes ONTAP is available in Google Cloud as a single- node system and as a high-availability (HA) pair of nodes. Based on the requirements, we can choose the Cloud Volumes ONTAP deployment mode. Upgrading a single node system to an HA pair is not supported. If you want to switch between a single- node system and an HA pair, then you must deploy a new system and replicate data from the existing system to the new system.

Highly available Cloud Volumes ONTAP in Google Cloud


Google Cloud supports deployment of resources across multiple geographical regions and multiple zones within a region. The HA deployment consists of two ONTAP nodes that use powerful n1-standard or n2-standard machine types available in Google Cloud. Data is synchronously replicated between the two Cloud Volumes ONTAP nodes to provide availability in the event of a failure. HA deployment of Cloud Volumes ONTAP requires four VPCs and a private subnet in each VPC. The subnets in the four VPCs should be provisioned with non-overlapping CIDR ranges.

The four VPCs are used for the following purposes:

- VPC 0 enables inbound communication to data and Cloud Volumes ONTAP nodes.
- VPC 1 provides cluster connectivity between Cloud Volumes ONTAP nodes.
- VPC 2 allows for non-volatile ram (NVRAM) replication between nodes.

- VPC 3 is used for connectivity to the HA mediator instance and disk replication traffic for node rebuilds.

The following image shows a highly available Cloud Volumes ONTAP in Goggle Cloud.



For details, see [this link](#).

For networking requirements for Cloud Volumes ONTAP in Google Cloud, see [this link](#).

For details about data tiering, see [this link](#).

Set up environment prerequisites

The automated creation of Cloud Volumes ONTAP clusters, SnapMirror configuration between an on-premises volume and a Cloud volume, creating a cloud volume, and so on are performed using Terraform configuration. These Terraform configurations are hosted on a Terraform Cloud for Business account. Using Intersight Cloud Orchestrator, you orchestrate tasks like creating a workspace in a Terraform Cloud for Business account, add all required variables to the workspace, execute a Terraform Plan, and so on.

For these automation and orchestration tasks, there are a few requirements and data needed, as is described in the following sections.

GitHub repository

You need a GitHub account to host your Terraform code. Intersight Orchestrator creates a new workspace in the Terraform Cloud for Business account. This workspace is configured with a version control workflow. For this purpose, you need to keep the Terraform configuration in a GitHub repository and provide it as an input while creating the workspace.

[This GitHub link](#) provides the Terraform configuration with various resources. You can fork this repository and make a copy in your GitHub account.

In this repository, `provider.tf` has the definition for the required Terraform provider. Terraform provider for NetApp Cloud Manager is used.

`variables.tf` has all the variable declarations. The value for these variables is input as the Intersight Cloud Orchestrator's workflow input. This provides a convenient way to pass values to a workspace and execute the Terraform configuration.

`resources.tf` defines the various resources needed to add an on-premises ONTAP to the working environment, create a single node Cloud Volumes ONTAP cluster on Google Cloud, establish a SnapMirror relationship between on-premises and Cloud Volumes ONTAP, create a cloud volume on Cloud Volumes ONTAP, and so on.

In this repository:

- `provider.tf` has NetApp Cloud Manager as a definition for the required Terraform provider.
- `variables.tf` has the variable declarations that are used as input for the Intersight Cloud Orchestrator workflow. This provides a convenient way to pass values to workspace and execute Terraform configuration.
- `resources.tf` defines various resources to add an on-premises ONTAP to the working environment, create a single- node Cloud Volumes ONTAP cluster on Google Cloud, establish a SnapMirror relationship between on-premises and Cloud Volumes ONTAP, create a cloud volume on Cloud Volumes ONTAP, and so on.

You can add an additional resource block to create multiple volumes on Cloud Volumes ONTAP or use `count` or `for_each` Terraform constructs.

To connect Terraform workspaces, modules, and policy sets to git repositories containing Terraform configurations, Terraform Cloud needs access to your GitHub repo.

Add a client, and the OAuth Token ID of the client is used as one of the Intersight Cloud Orchestrator's workflow input.

1. Log in to your Terraform Cloud for Business account. Navigate to **Settings > Providers**.
2. Click **Add a VCS provider**.
3. Select your version.
4. Follow the steps under **Set up provider**.
5. You see the added client in **VCS Providers**. Make a note of the OAuth Token ID.

Refresh token for NetApp Cloud Manager API operations

In addition to the web browser interface, Cloud Manager has a REST API that provides software developers with direct access to the Cloud Manager functionality through the SaaS interface. The Cloud Manager service

consists of several distinct components that collectively form an extensible development platform. The refresh token enables you to generate access tokens that you add to the Authorization header for each API call.

Without calling an API directly, the netapp-cloudmanager provider uses a refresh token and translates the Terraform resources into corresponding API calls. You need to generate a refresh token for NetApp Cloud Manager API operations from [NetApp Cloud Central](#).

You need the client ID of the Cloud Manager Connector to create resources on Cloud Manager such as creating a Cloud Volumes ONTAP cluster, configuring SnapMirror, and so on.

1. Log into Cloud Manager: <https://cloudmanager.netapp.com/>.
2. Click **Connector**.
3. Click **Manage Connectors**.
4. Click the ellipses and copy the Connector ID.

Develop Cisco Intersight Cloud Orchestrator workflow

Cisco Intersight Cloud Orchestrator is available in Cisco Intersight if:

- You have installed the Intersight Premier license.
- You are either an account administrator, storage administrator, virtualization administrator, or server administrator and have a minimum of one server assigned to you.

Workflow Designer

The Workflow Designer helps you create new workflows (as well as tasks and data types) and edit existing workflows to manage targets in Cisco Intersight.

To launch the Workflow Designer, go to **Orchestration > Workflows**. A dashboard displays the following details under the tabs **My Workflows**, **Sample Workflows**, and **All Workflows**:

- Validation Status
- Last Execution Status
- Top Workflows by Execution Count
- Top Workflow Categories
- Number of System Defined Workflows
- Top Workflows by Targets

Using the dashboard, you can create, edit, clone, or delete a tab. To create your own custom view tab, click **+**, specify a name, and then select the required parameters that need to be displayed in the columns, tag columns, and widgets. You can rename a tab if it doesn't have a **Lock** icon.

Under the dashboard is a tabular list of workflows displaying the following information:

- Display Name
- Description
- System Defined
- Default Version
- Executions

- Last Execution Status
- Validation Status
- Last Update
- Organization

The Actions column allows you to perform the following actions:

- **Execute.** Executes the workflow.
- **History.** Displays workflow execution history.
- **Manage Versions.** Create and manage versions for workflows.
- **Delete.** Delete a workflow.
- **Retry.** Retry a failed workflow.

Workflow

Create a workflow that consists of the following steps:

- **Defining a workflow.** Specify the display name, description, and other important attributes.
- **Define workflow inputs and workflow outputs.** Specify which input parameters are mandatory for the workflow execution, and the outputs generated on successful execution
- **Add workflow tasks.** Add one or more workflow tasks in the Workflow Designer that are needed for the workflow to carry out its function.
- ***Validate the workflow.** *Validate a workflow to ensure that there are no errors in connecting task inputs and outputs.

Create workflows for on-premises FlexPod storage

To configure a workflow for on-premises FlexPod storage, see [this link](#).

Next: [DR workflow](#).

DR workflow

Previous: [Automated deployment of Hybrid Cloud NetApp Storage](#).

The sequence of steps are as follows:

1. Define the workflow.
 - Create a short, user-friendly name for the workflow, such as Disaster Recovery Workflow.
2. Define the workflow input. The inputs we take for this workflow include the following:
 - Volume options (volume name, mount path)
 - Volume capacity
 - Data center associated with the new datastore
 - Cluster on which the datastore is hosted
 - Name for the new datastore to create in vCenter
 - Type and version of the new datastore

- Name of the Terraform organization
- Terraform workspace
- Description of the Terraform workspace
- Variables (sensitive and nonsensitive) required to execute Terraform configuration
- Reason for starting the plan

3. Add the workflow tasks.

The tasks related to operations in FlexPod include the following:

- Create volume in FlexPod.
- Add storage export policy to the created volume.
- Map the newly created volume to a datastore in VMware vCenter.

The tasks related to creating Cloud Volumes ONTAP cluster:

- Add Terraform workspace
- Add Terraform variables
- Add Terraform sensitive variables
- Start new Terraform plan
- Confirm Terraform run

4. Validate the workflow.


Procedure 1: Create the workflow

1. Click **Orchestration** from the left navigation pane and click **Create Workflow**.
2. In the **General** tab:
 - a. Provide the display name (Disaster Recovery Workflow).
 - b. Select the organization, set tags, and provide a description.
3. Click Save.


The screenshot shows the 'General' tab of a workflow configuration interface. The 'Display Name' field contains 'Disaster Recovery Workflow'. The 'Reference Name' field contains 'DisasterRecoveryWorkflow'. The 'Organization' dropdown is set to 'default'. The 'Version' dropdown is set to '2 (default)'. The 'Description' field contains 'Workflow which creates and configures SnapMirror between FlexPod Storage and Cloud Volumes ONTAP'. Under the 'Workflow Execution' section, the 'Failed/Terminated Actions' checkbox is checked. The 'Enable Retry' checkbox is checked, 'Enable Auto Rollback' is unchecked, and 'Enable Debug Logs' is checked. At the bottom, there are three tabs: 'Workflow Inputs', 'Workflow Variables', and 'Workflow Outputs'. The 'Workflow Inputs' tab is selected, and an 'Add Workflow Input' button is visible below it.

Procedure 2. Create a new volume in FlexPod

1. Go to the **Designer** tab and click **Tasks** from the **Tools** section.
2. Drag and drop the **Storage > New Storage Volume** task from the **Tools** section into the **Design** area.
3. Click **New Storage Volume**.

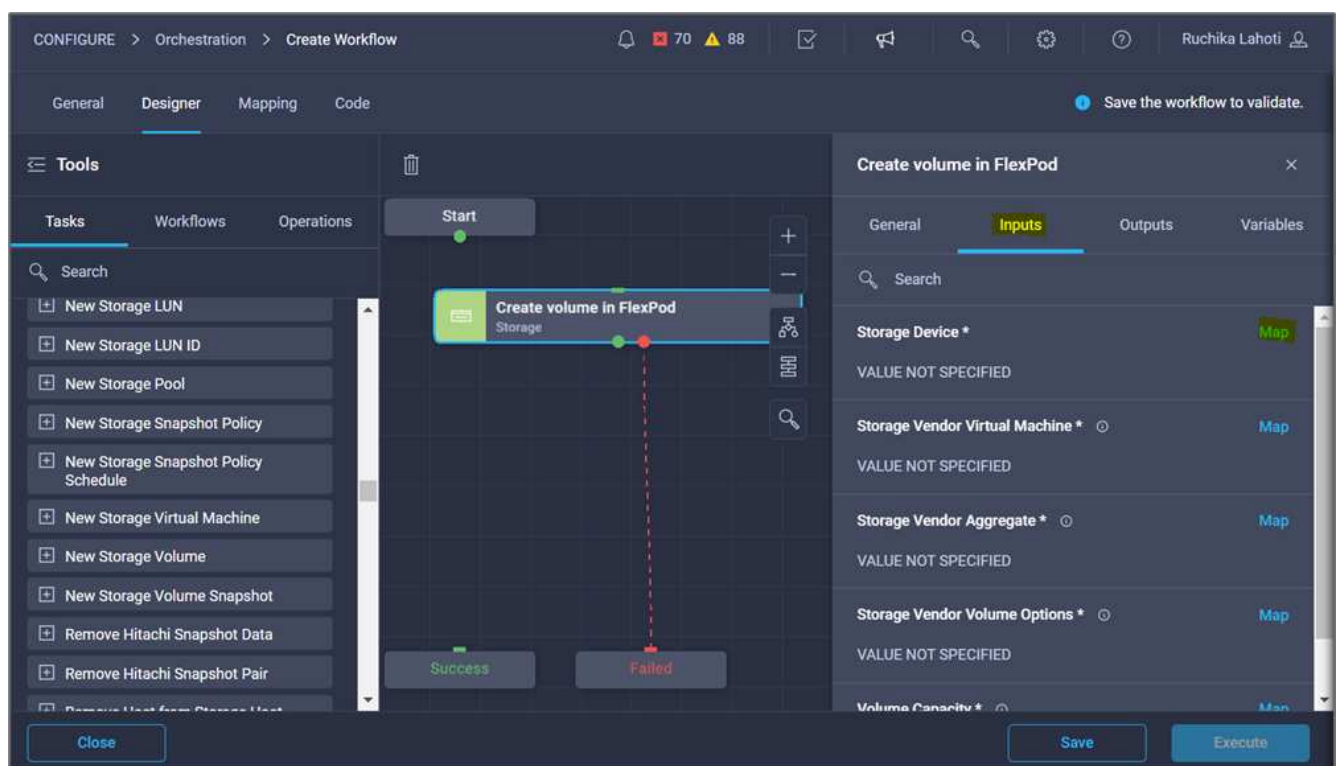


4. In the **Task Properties** area, click the **General** tab. Optionally, you can change the name and description for this task. In this example, the name of the task is **Create Volume in FlexPod**.




5. In the **Task Properties** area, click **Inputs**.

6. Click **Map** in the **Storage Device** field.




7. Choose **Static Value** and click **Select Storage Device**.


8. Click the storage target added and click **Select**.




9. Click **Map**.




10. Click **Map** in the **Storage Vendor Virtual Machine** field.




11. Choose **Static Value** and click **Select Storage Virtual Machine**.




12. Select the storage virtual machine where the volume needs to be created and click **Select**.



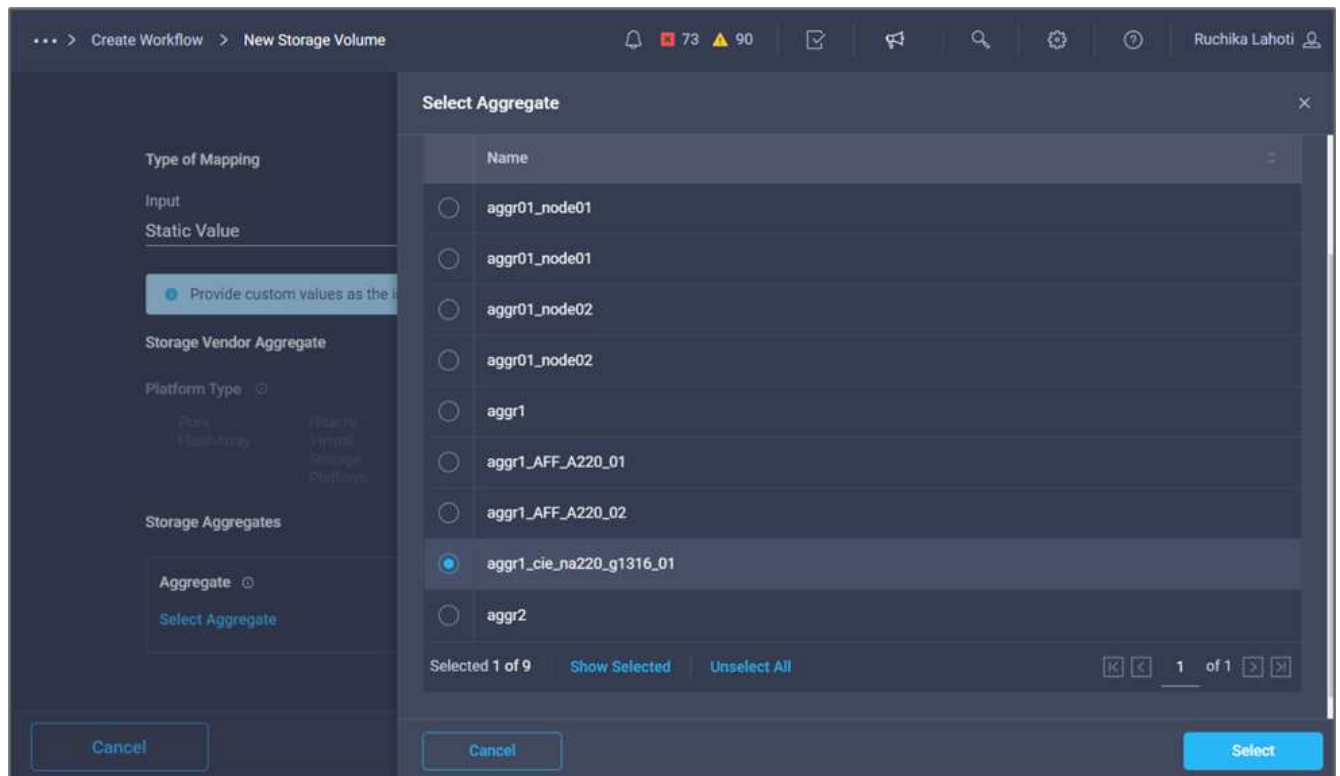
13. Click **Map**.



14. Click **Map** in the **Storage Vendor Aggregate** field.



15. Choose **Static Value** and click **Select Storage Aggregate**. Choose the aggregate and click **Select**.



16. Click **Map**.
17. Click **Map** in the **Storage Vendor Volume Options** field.
18. Choose **Direct Mapping** and click **Workflow Input**.

... > Create Workflow > New Storage Volume

73 90

Map Task Input

Configure/Assign the value from available options.

Type of Mapping

Input

Direct Mapping v ⓘ

Map the workflow input, variable or any of the previous task's outputs to input.

Map to

Workflow Input v ⓘ

Input Name * v ⓘ

Add Workflow Input

19. In the Add Input wizard, complete the following steps:
- Provide a display name and reference name (optional).
 - Make sure that **Storage Vendor Volume Options** is selected for the **Type**.
 - Click **Set Default Value and Override**.
 - Click **Required**.
 - Set the **Platform Type** to **NetApp Active IQ Unified Manager**.
 - Provide a default value for the created volume under **Volume**.
 - Click **NFS**. If NFS is set, an NFS volume is created. If this value is set to false, a SAN volume is created.
 - Provide a mount path and click **Add**.

Add Workflow Input

☒ Set Default Value ⓘ

☒ Allow User Override ⓘ

Default Values *

Storage Vendor Volume Options

Platform Type ⓘ

☐ Pure FlashArray ☐ Hitachi Virtual Storage Platform ☒ NetApp Active IQ Unified Manager ☐ None

Volume *

mssql_data_vol ⓘ

NFS Volume Option

☒ NFS ⓘ

Mount Path

/mssql_data_vol ⓘ

Cancel Add

20. Click **Map**.
21. Click **Map** in the **Volume Capacity** field.
22. Choose **Direct Mapping** and click **Workflow Input**.
23. Click **Input Name** and **Create Workflow Input**.

... > Create Workflow > New Storage Volume > Volume Capacity

73 90

Ruchika Lahoti

Map Task Input

Configure/Assign the value from available options.

Type of Mapping

Input

Direct Mapping

Map the workflow input, variable or any of the previous task's outputs to input.

Map to

Workflow Input

Input Name *


Add Workflow Input

Storage Vendor Volume Options


Cancel

Map

24. In the Add Input wizard:
- Provide a display name and a reference name (optional).
 - Click **Required**.
 - For **Type**, select **Storage Capacity**.
 - Click **Set Default Value and Override**.
 - Provide a default value for the volume size and unit.
 - Click **Add**.



25. Click **Map**.
26. With Connector, create a connection between the **Start** and **Create Volume in FlexPod** tasks, and click **Save**.






Ignore the error for now. This error displays because there is no connectivity between the tasks **Create Volume in FlexPod** and **Success** which is required to specify the successful transition.

Procedure 3: Add storage export policy


1. Go to the **Designer** tab and click **Tasks** from the **Tools** section.
2. Drag and drop the **Storage > Add Storage Export Policy to Volume** task from the **Tools** section in the **Design** area.
3. Click **Add Storage Export Policy to Volume**. In the **Task Properties** area, click the **General** tab. Optionally, you can change the name and description for this task. In this example, the name of the task is Add Storage Export Policy.
4. Use Connector to make a connection between the tasks **Create Volume in FlexPod** and **Add Storage Export Policy**. Click **Save**.

The screenshot shows the NetBackup Orchestrator Designer interface. The top navigation bar includes 'CONFIGURE > Orchestration > Disaster recovery workflow > Edit'. The left sidebar shows the 'Tools' section with 'Tasks' selected. The main canvas displays a workflow starting with 'Start', followed by 'Create volume in FlexPod' (Storage task), then 'Add Storage Export Policy to V...' (Storage task), and finally branching to 'Success' and 'Failed' endpoints. The 'Task Properties' panel on the right is open for the 'Add Storage Export Policy to Volume' task, showing the 'General' tab with fields for Name, Version, Task Type, and User Description.


5. In the **Task Properties** area, click **Inputs**.
6. Click **Map** in the **Storage Device** field.



7. Choose **Static Value** and click **Select Storage Device**. Select the same storage target added while creating the previous task of creating a new storage volume.
8. Click **Map**.



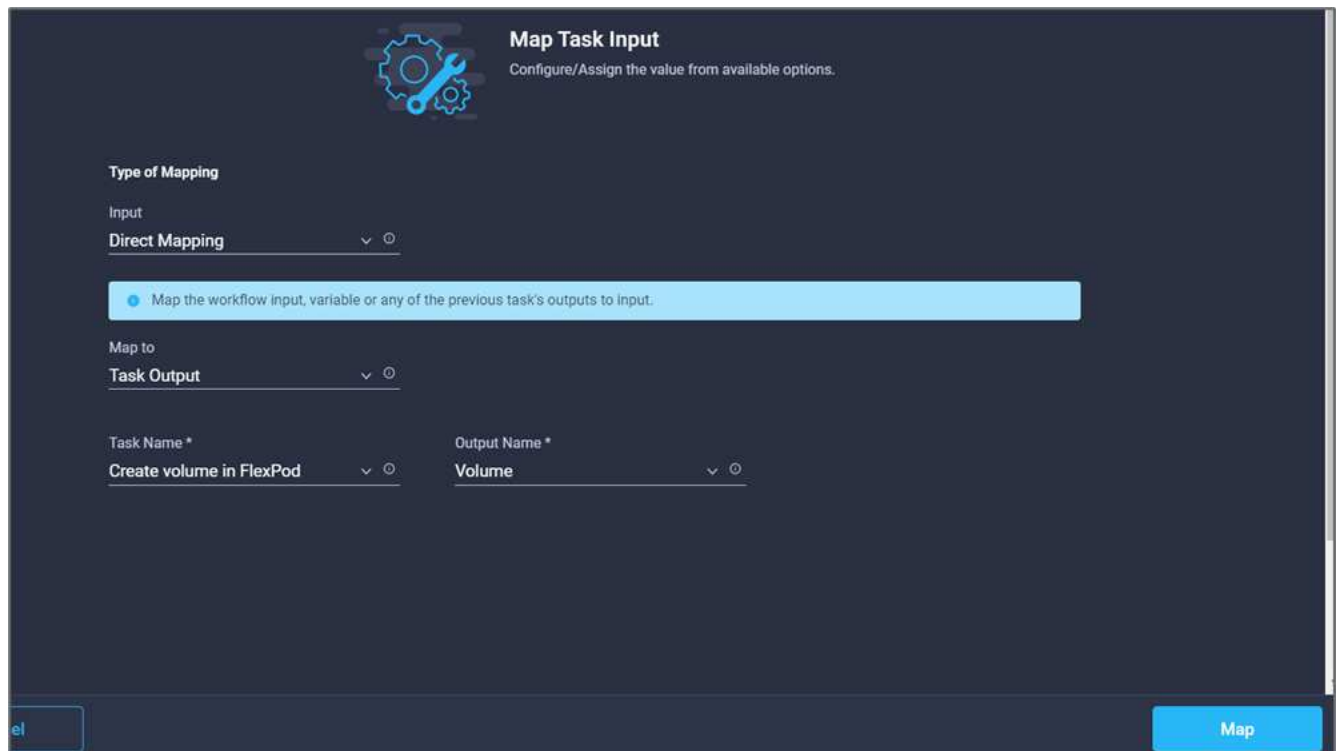
9. Click **Map** in the **Storage Vendor Virtual Machine** field.
10. Choose **Static Value** and click **Select Storage Virtual Machine**. Select the same storage virtual machine added while creating the previous task of creating a new storage volume.




11. Click **Map**.
12. Click **Map** in the **Volume** field.
13. Click **Task Name** and then click **Create Volume in FlexPod**. Click **Output Name** and then **Volume**.



In Cisco Intersight Cloud Orchestrator, you can provide the output of a previous task as the input for a new task. In this example, the **Volume** details were provided from the **Create Volume in FlexPod** task as an input for the task **Add Storage Export Policy**.



14. Click **Map**.
15. Click **Map** in the **Export Policy** field.
16. Choose **Static Value** and click **Select Export Policy**. Select the export policy created.



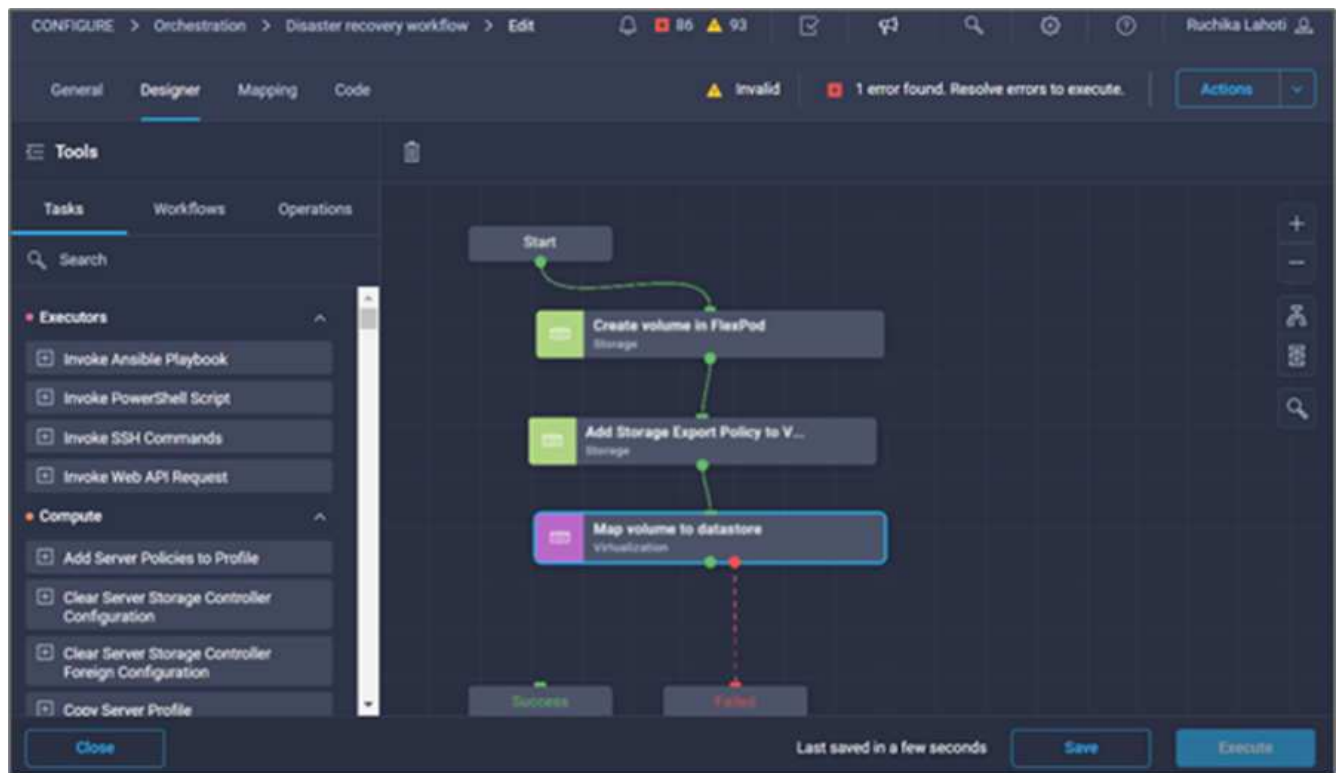
17. Click **Map** and then **Save**.




This completes addition of an export policy to the volume. Next, you create a new datastore mapping the created volume.

Procedure 4: Map FlexPod volume to datastore


1. Go to the **Designer** tab and click **Tasks** from the **Tools** section.
2. Drag and drop the **Virtualization > New Hypervisor Datastore** task from the **Tools** section in the **Design** area.
3. Use Connector to make a connection between the **Add Storage Export Policy** and **New Hypervisor Datastore** tasks. Click **Save**.



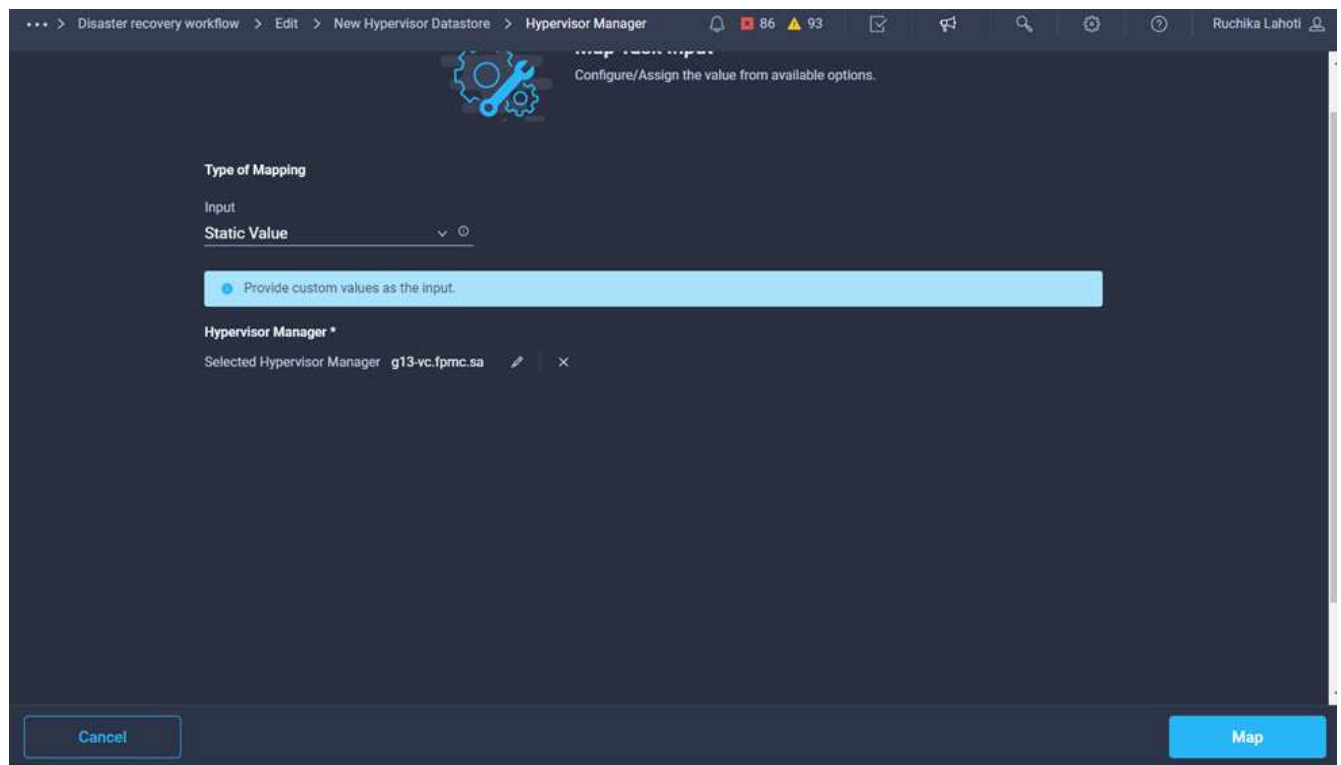
4. Click **New Hypervisor Datastore**. In the **Task Properties** area, click the **General** tab. Optionally, you can change the name and description for this task. In this example, the name of the task is **Map volume to Datastore**.



5. In the **Task Properties** area, click **Inputs**.
6. Click **Map** in the **Hypervisor Manager** field.
7. Choose **Static Value** and click **Select Hypervisor Manager**. Click the VMware vCenter target.



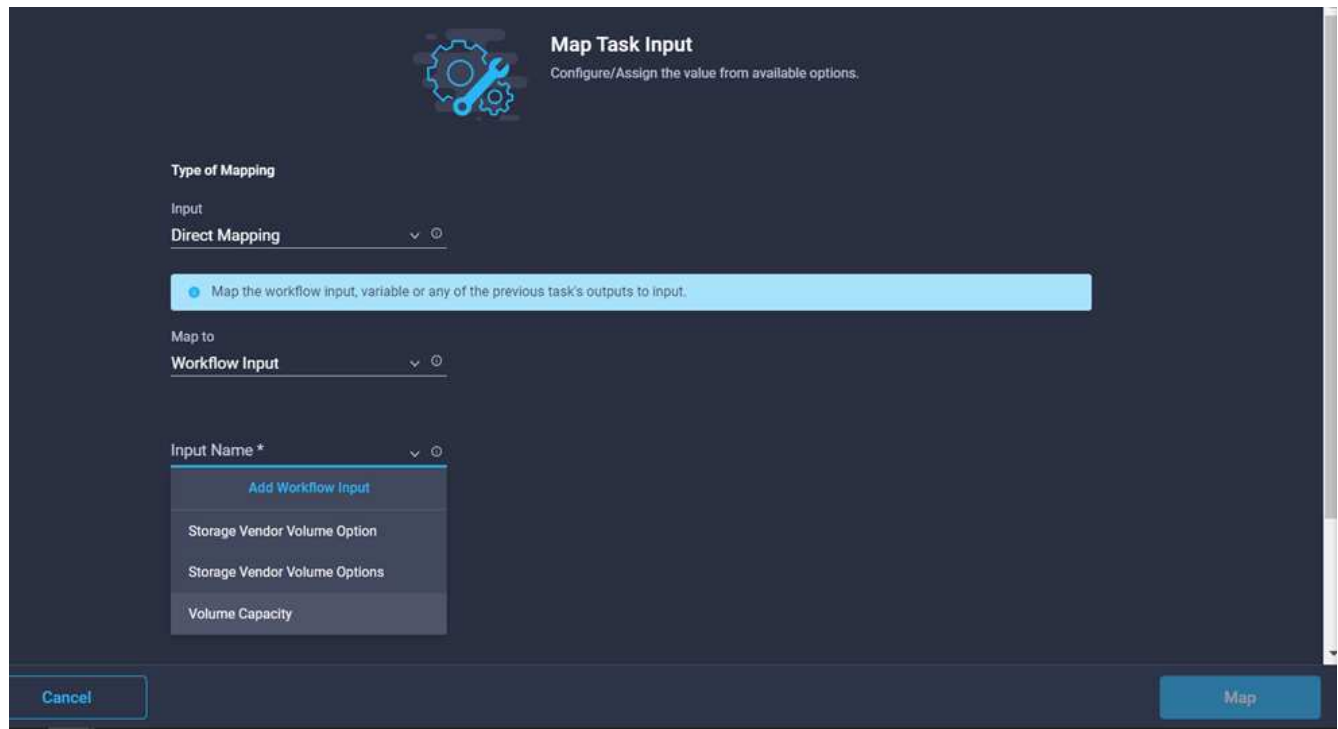
8. Click **Map**.



9. Click **Map** in the **Data center** field. This is the data center associated with the new datastore.

10. Choose **Direct Mapping** and click **Workflow Input**.

11. Click **Input Name** and then **Create Workflow Input**.



Map Task Input
Configure/Assign the value from available options.

Type of Mapping
Input
Direct Mapping

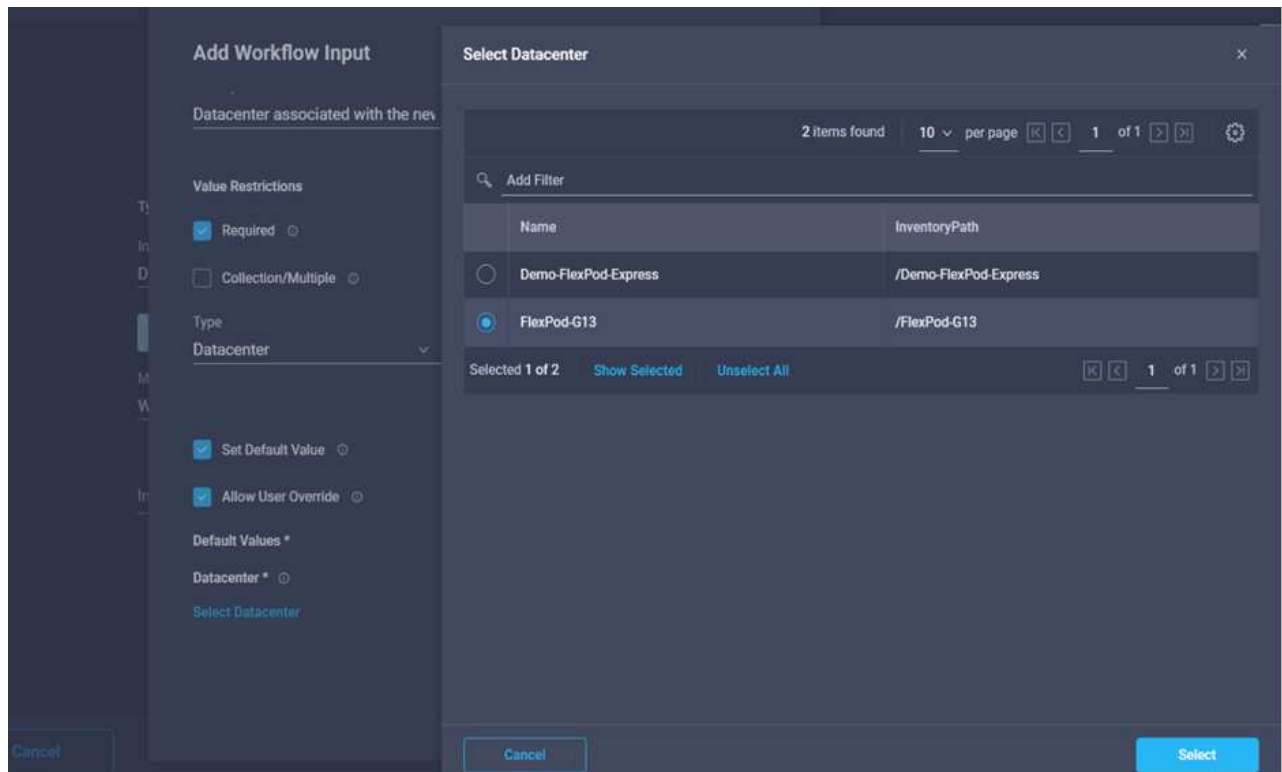
Map the workflow input, variable or any of the previous task's outputs to input.

Map to
Workflow Input

Input Name *
Add Workflow Input
Storage Vendor Volume Option
Storage Vendor Volume Options
Volume Capacity

Cancel Map

12. In the Add Input wizard, complete the following steps:
 - a. Provide a display name and reference name (optional).
 - b. Select **Datacenter** as the type.
 - c. Click **Set Default Value and Override**.
 - d. Click **Select Datacenter**.
 - e. Click the data center associated with the new datastore and then click **Select**.



Add Workflow Input

Datacenter associated with the new workflow input

Value Restrictions
☒ Required
☐ Collection/Multiple

Type
Datacenter

☒ Set Default Value
☒ Allow User Override

Default Values *
 Datacenter *
[Select Datacenter](#)

Select Datacenter

2 items found | 10 per page | 1 of 1

Name	InventoryPath
Demo-FlexPod-Express	/Demo-FlexPod-Express
FlexPod-G13	/FlexPod-G13

Selected 1 of 2 | [Show Selected](#) | [Unselect All](#) | 1 of 1

Cancel Select

- Click **Add**.

13. Click **Map**.

14. Click **Map** in the **Cluster** field.

15. Choose **Direct Mapping** and click **Workflow Input**.

Map Task Input
Configure/Assign the value from available options.

Type of Mapping
Input
Direct Mapping

Map the workflow input, variable or any of the previous task's outputs to input.


Map to
Workflow Input

Input Name *
Add Workflow Input
Datacenter
Storage Vendor Volume Option
Storage Vendor Volume Options
Volume Capacity

Cancel Map

16. In the Add Input wizard, complete the following steps:

- Provide a display name and reference name (optional).
- Click **Required**.
- Select Cluster as the type.
- Click **Set Default Value and Override**.
- Click **Select Cluster**.
- Click the cluster associated with the new datastore.
- Click **Select**.



h. Click **Add**.

17. Click **Map**.

18. Click **Map** in the **Host** field.

Add Workflow Input

Cluster on which the datastore will ⓘ

Value Restrictions

☐ Required ⓘ

☐ Collection/Multiple ⓘ

Type

Cluster ▼ ⓘ

☒ Set Default Value ⓘ

☒ Allow User Override ⓘ

Default Values *

Cluster ⓘ

Selected Cluster Washington ✎ ✕

Cancel Add

19. Choose **Static Value** and click the host on which the datastore will be hosted. If a cluster is specified, then the host is ignored.

4 items found | 10 per page | 1 of 1

Add Filter

	Name
<input checked="" type="radio"/>	172.22.0.111
<input type="radio"/>	172.22.0.112
<input type="radio"/>	esxi-01.nva.local
<input type="radio"/>	esxi-02.nva.local

Selected 1 of 4 | Show Selected | Unselect All | 1 of 1

Cancel | Select

20. Click **Select and Map**.
21. Click **Map** in the **Datastore** field.
22. Choose **Direct Mapping** and click **Workflow Input**.
23. Click **Input Name** and **Create Workflow Input**.

Map Task Input

Configure/Assign the value from available options.

Type of Mapping

Input

Direct Mapping

Map the workflow input, variable or any of the previous task's outputs to input.

Map to

Workflow Input

Input Name *

- Add Workflow Input
- Cluster
- Datacenter
- Storage Vendor Volume Option
- Storage Vendor Volume Options
- Volume Capacity

Cancel

24. In the Add Input wizard:
 - a. Provide a display name and reference name (optional).
 - b. Click **Required**.
 - c. Click **Set Default Value and Override**.
 - d. Provide a default value for the datastore and click **Add**.

Add Workflow Input

Type
String

Min 0 Max 0 Regex ^.{1,42}\$

☐ Secure

☐ Object Selector

☒ Set Default Value

☒ Allow User Override

Default Values *

Datastore *
hybrid-ds

Cancel Add

25. Click **Map**.
26. Click **Map** in the input field **Type of Datastore**.
27. Choose **Direct Mapping** and click **Workflow Input**.
28. Click **Input Name** and **Create Workflow Input**.

Type of Mapping

Input

Direct Mapping ▼ ⓘ

• Map the workflow input, variable or any of the previous task's outputs to input.

Map to

Workflow Input ▼ ⓘ

Input Name * ▼ ⓘ

- Add Workflow Input
- Cluster
- Datacenter
- Datastore
- Storage Vendor Volume Option
- Storage Vendor Volume Options

Map

29. In the Add Input wizard, complete the following steps:
- Provide a display name and reference name (optional) and click **Required**.
 - Make sure to select the type **Types of Datastore** and click **Set Default Value and Override**.

Add Workflow Input

Display Name *
Type of Datastore

Reference Name *
DatastoreVersion

Description
Type and version of the new datast

Value Restrictions

☒ Required

☐ Collection/Multiple

Type
Types of Datastore


☒ Set Default Value

☒ Allow User Override

Default Values *
Type of Datastore

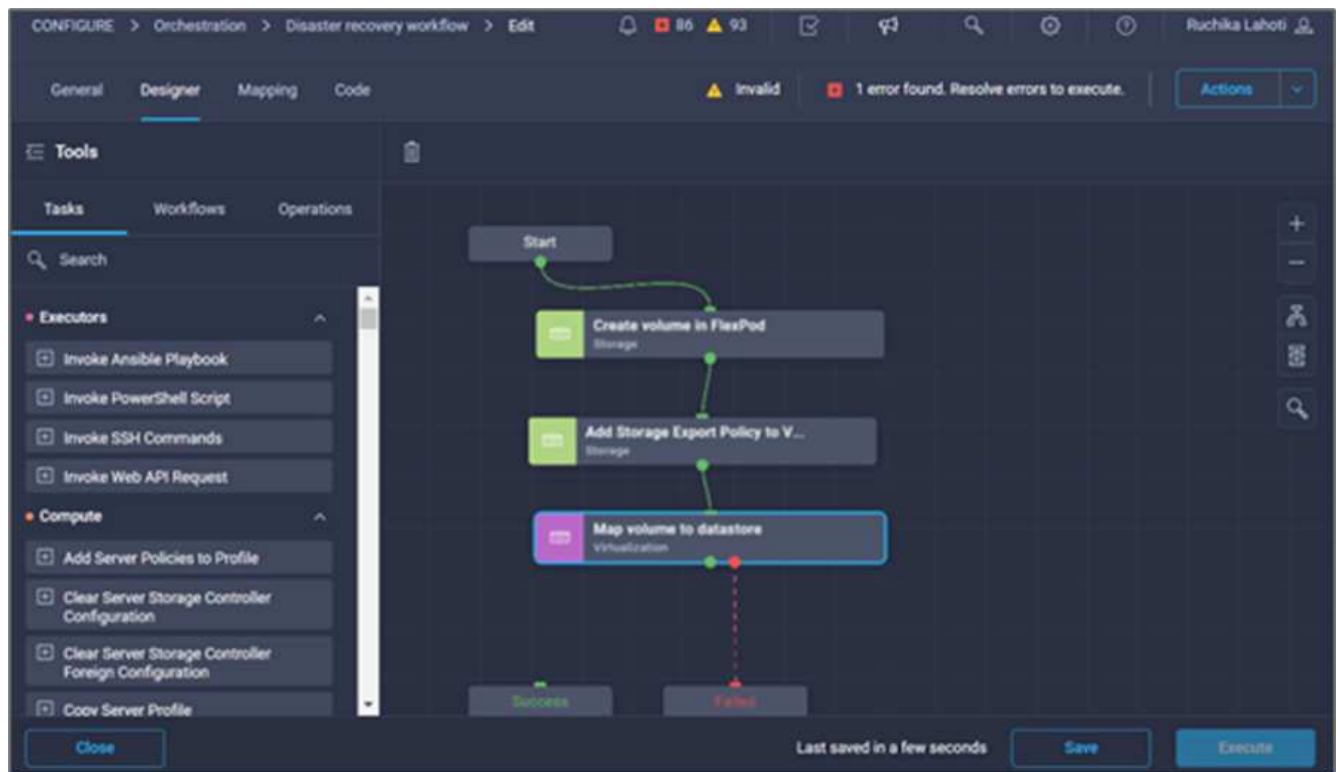
Cancel Add

- c. Provide the Remote Path. This is the remote path of the NFS mount point.
- d. Provide the host names or IP addresses of remote NFS server in NFS Server Address.
- e. Click the **Access Mode**. The Access mode is for the NFS server. Click read-only if volumes are exported as read-only. Click **Add**.




30. Click **Map**.

31. Click **Save**.

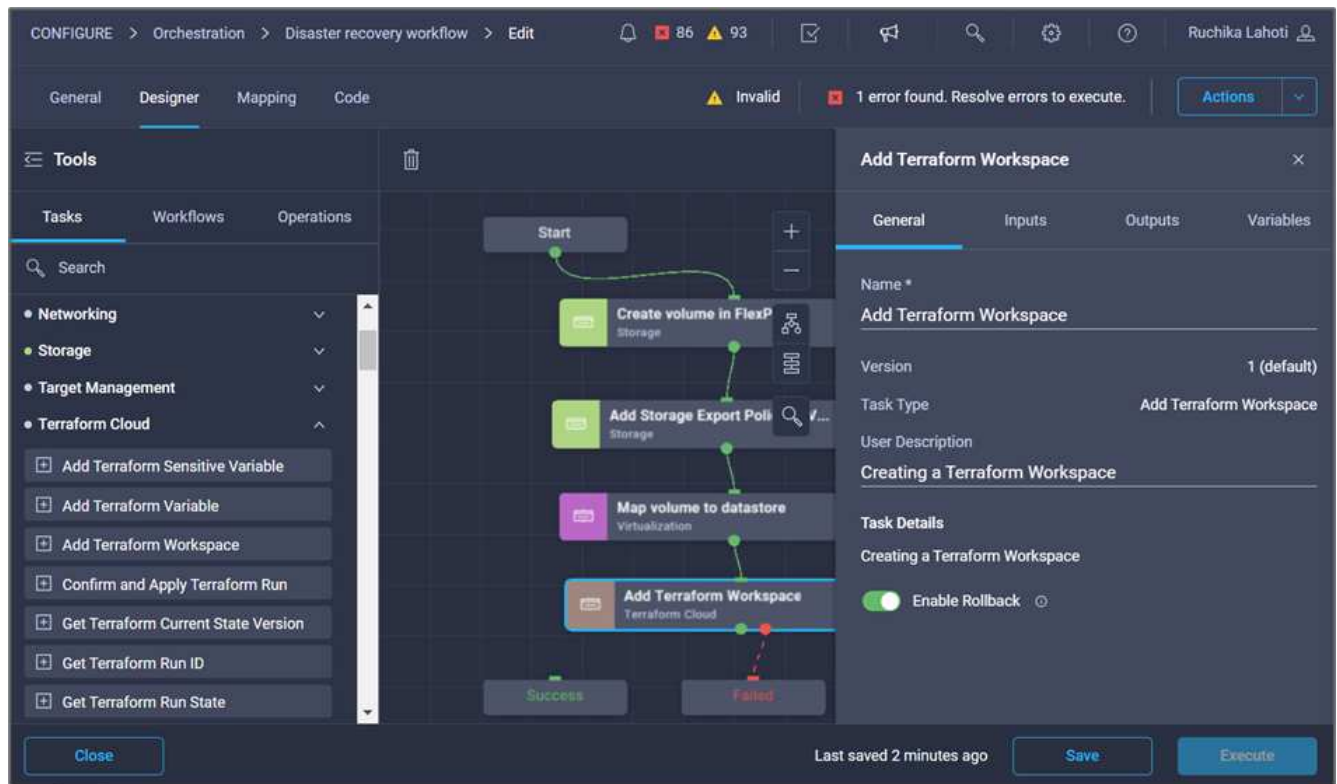


This completes the task of creating the datastore. All the tasks performed in the on- premises FlexPod Datacenter are completed.




Procedure 5: Add a new Terraform workspace

1. Go to the **Designer** tab and click **Tasks** from the **Tools** section.
2. Drag and drop the **Terraform Cloud > Add Terraform Workspace** task from the Tools section in the Design area.
3. Use Connector to connect the **Map volume to Datastore** and **Add Terraform Workspace** tasks and click **Save**.
4. Click **Add Terraform Workspace**. In the Task Properties area, click the **General** tab. Optionally, you can change the Name and Description for this task.




5. In the Task Properties area, click **Inputs**.
6. Click **Map** in the input field **Terraform Cloud Target**.
7. Choose **Static Value** and click **Select Terraform Cloud Target**. Select the Terraform Cloud for Business account that was added as explained in [Configure Cisco Intersight Service for HashiCorp Terraform.](#)

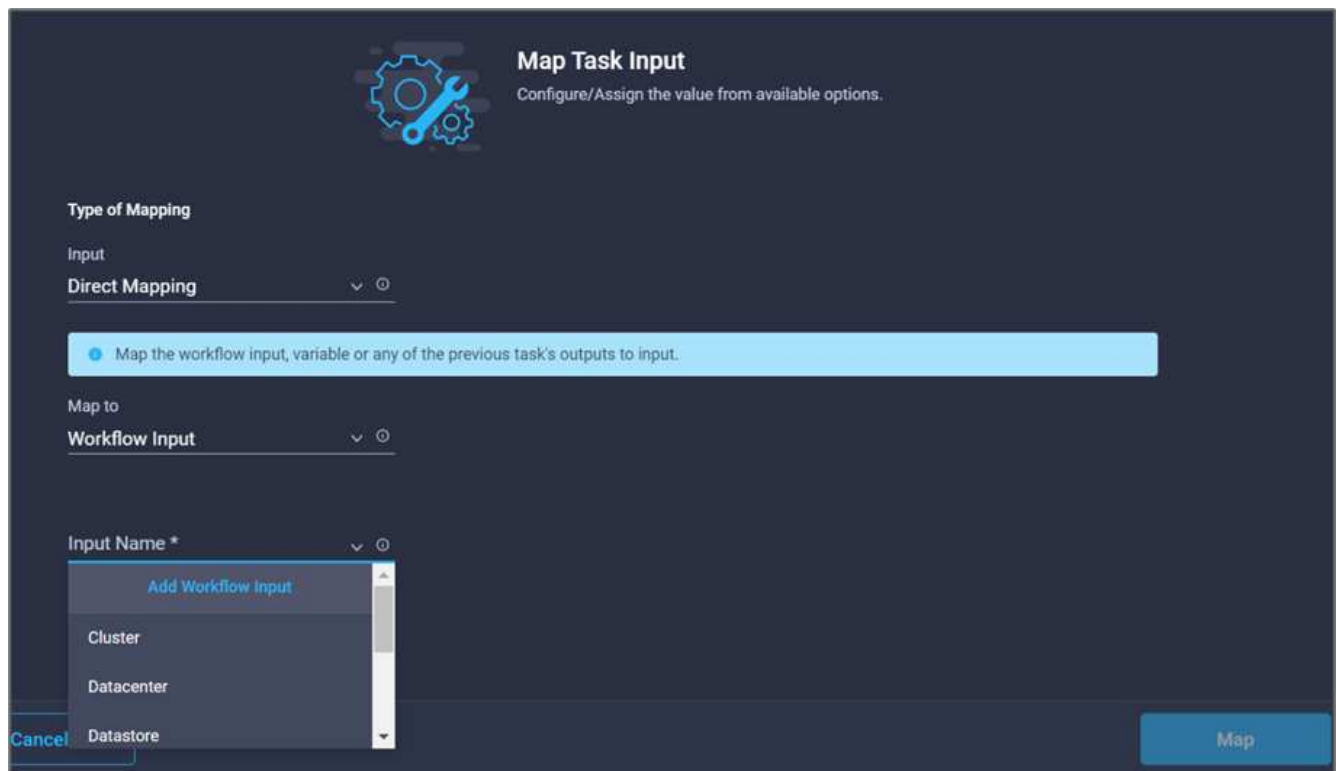


8. Click **Map**.

9. Click **Map** in the input field **Terraform Organization Name**.
10. Choose **Static Value** and then click **Select Terraform Organization**. Select the name of the Terraform Organization that you are part of in your Terraform Cloud for Business account.



11. Click **Map**.
12. Click **Map** in the **Terraform Workspace Name** field. This is the new workspace in the Terraform Cloud for Business account.
13. Choose **Direct Mapping** and click **Workflow Input**.
14. Click **Input Name** and **Create Workflow Input**.



The image shows a 'Map Task Input' dialog box with a dark blue background. At the top left is a gear and wrench icon. The title 'Map Task Input' is at the top right, with a subtitle 'Configure/Assign the value from available options.' below it. The 'Type of Mapping' section has a dropdown menu set to 'Input', and the 'Direct Mapping' option is selected. A light blue instruction bar says 'Map the workflow input, variable or any of the previous task's outputs to input.' The 'Map to' section has a dropdown menu set to 'Workflow Input'. The 'Input Name *' section has a dropdown menu with 'Add Workflow Input' at the top, followed by 'Cluster', 'Datacenter', and 'Datastore'. A 'Cancel' button is on the bottom left, and a 'Map' button is on the bottom right.

Map Task Input
Configure/Assign the value from available options.

Type of Mapping
Input
Direct Mapping


Map the workflow input, variable or any of the previous task's outputs to input.

Map to
Workflow Input

Input Name *
Add Workflow Input
Cluster
Datacenter
Datastore

Cancel Map

15. In the Add Input wizard, complete the following steps:
 - a. Provide a display name and reference name (optional).
 - b. Click **Required**.
 - c. Make sure to select **String** for **Type**.
 - d. Click **Set Default Value and Override**.
 - e. Provide a default name for workspace.
 - f. Click **Add**.



16. Click **Map**.
17. Click **Map** in the **Workspace Description** field.
18. Choose **Direct Mapping** and click **Workflow Input**.
19. Click **Input Name** and **Create Workflow Input**.

Add Workflow Input ✕

Workspace Description ⓘ WorkspaceDescription ⓘ

Description
Description of the Terraform Work: ⓘ

Value Restrictions

☐ Required ⓘ

☐ Collection/Multiple ⓘ

Type
String ▼ ⓘ

Min 0 ⓘ Max 0 ⓘ Regex ⓘ

☐ Secure ⓘ

☒ Object Selector ⓘ

☒ Set Default Value ⓘ

☒ Allow User Override ⓘ

Cancel Add

20. In the Add Input wizard, complete the following steps:
- Provide a display name and reference name (optional).
 - Make sure to select **String** for **Type**.
 - Click **Set Default Value and Override**.
 - Provide a workspace description and click **Add**.

Add Workflow Input

Value Restrictions

☐ Required ⓘ

☐ Collection/Multiple ⓘ

Type
String ▼ ⓘ

Min **0** ⓘ Max **0** ⓘ Regex ⓘ

☐ Secure ⓘ

☐ Object Selector ⓘ

☒ Set Default Value ⓘ

☒ Allow User Override ⓘ

Default Values *

Workspace Description
workspace to create CVO and configure SnapMirror ⓘ

Cancel Add

21. Click **Map**.
22. Click **Map** in the **Execution Mode** field.
23. Choose **Static Value**, click **Execution Mode**, and then click **remote**.

Type of Mapping

Input
 Static Value

Provide custom values as the input.

Execution Mode

ExecutionMode
 remote

24. Click **Map**.
25. Click **Map** in the **Apply Method** field.
26. Choose **Static Value** and click **Apply Method**. Click **Manual Apply**.

Type of Mapping


Input
 Static Value

Provide custom values as the input.

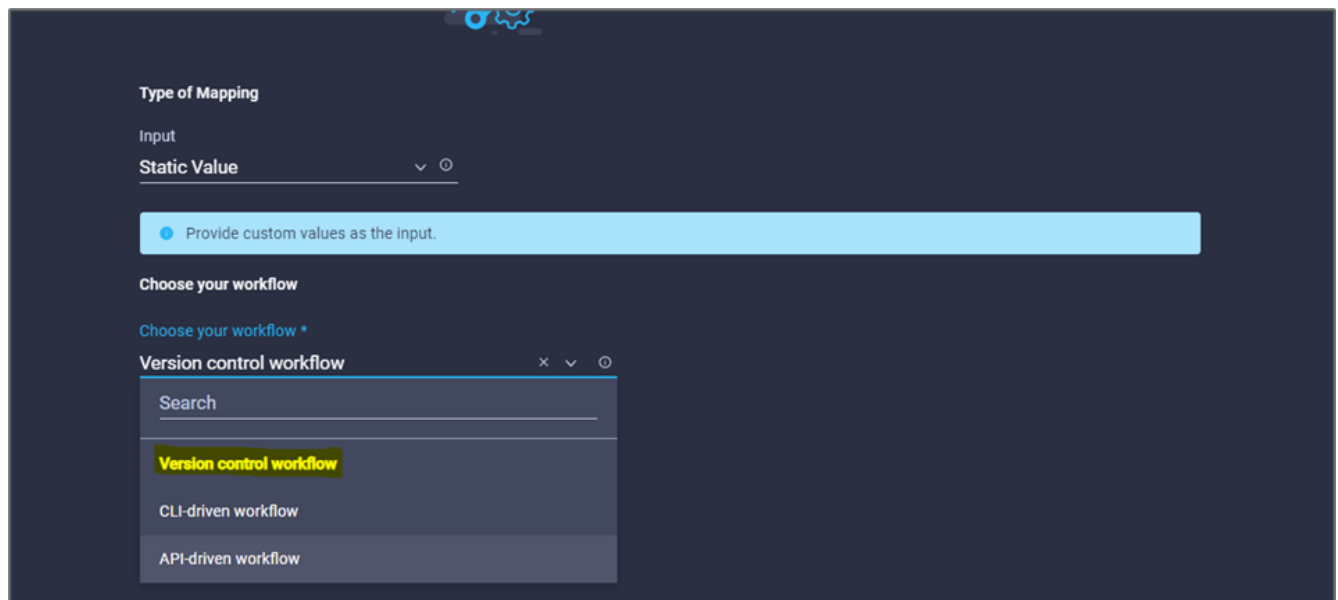
Apply Method

Manual Apply

27. Click **Map**.
28. Click **Map** in the **User Interface** field.
29. Choose **Static Value** and click **User Interface**. Click **Console UI**.

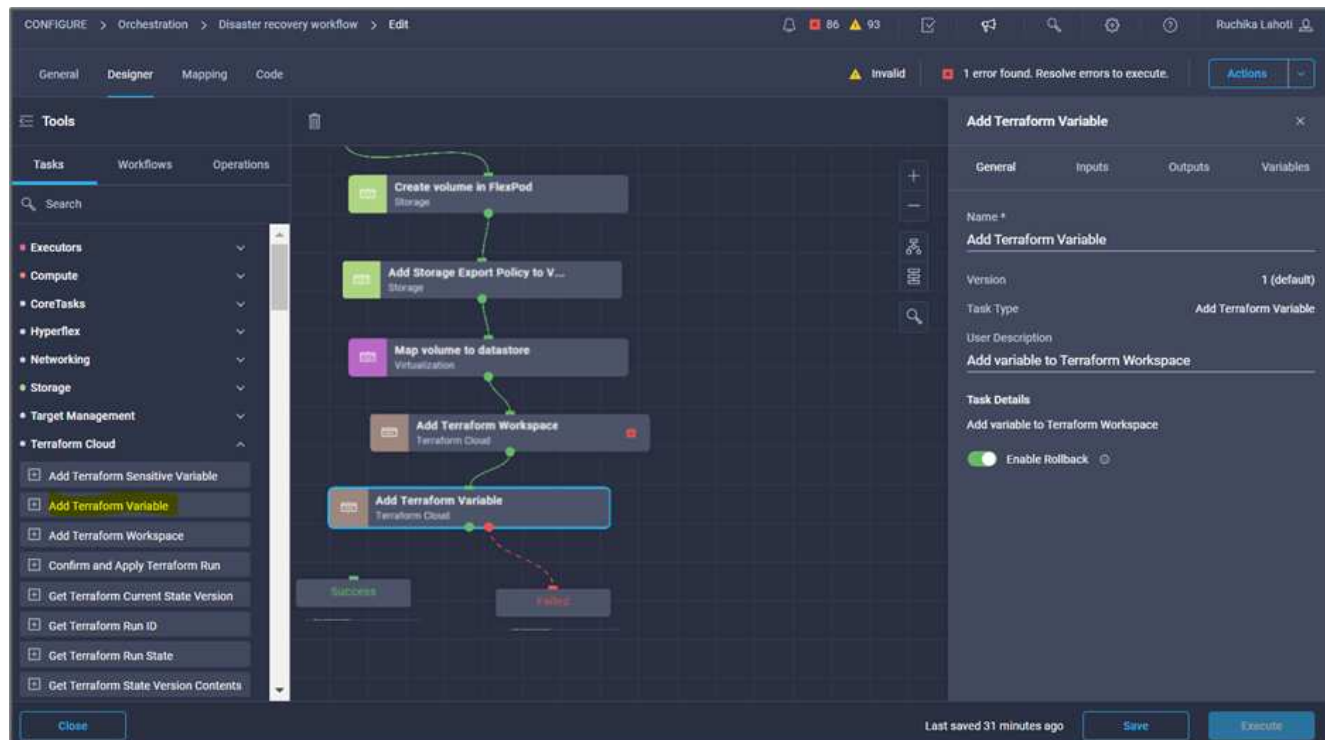


30. Click **Map**.
31. Click **Map** in the input field and select your workflow.
32. Select **Static Value**, and click **Choose Your Workflow**. Click **Version Control Workflow**.




33. Provide the following GitHub repository details:
 - a. In **Repository Name**, enter the name of the repository detailed in the section [“Set up environment prerequisites”](#).
 - b. Provide the OAuth Token ID as detailed in the section [“Set up environment prerequisites”](#).
 - c. Select the **Automatic Run Triggering** option.

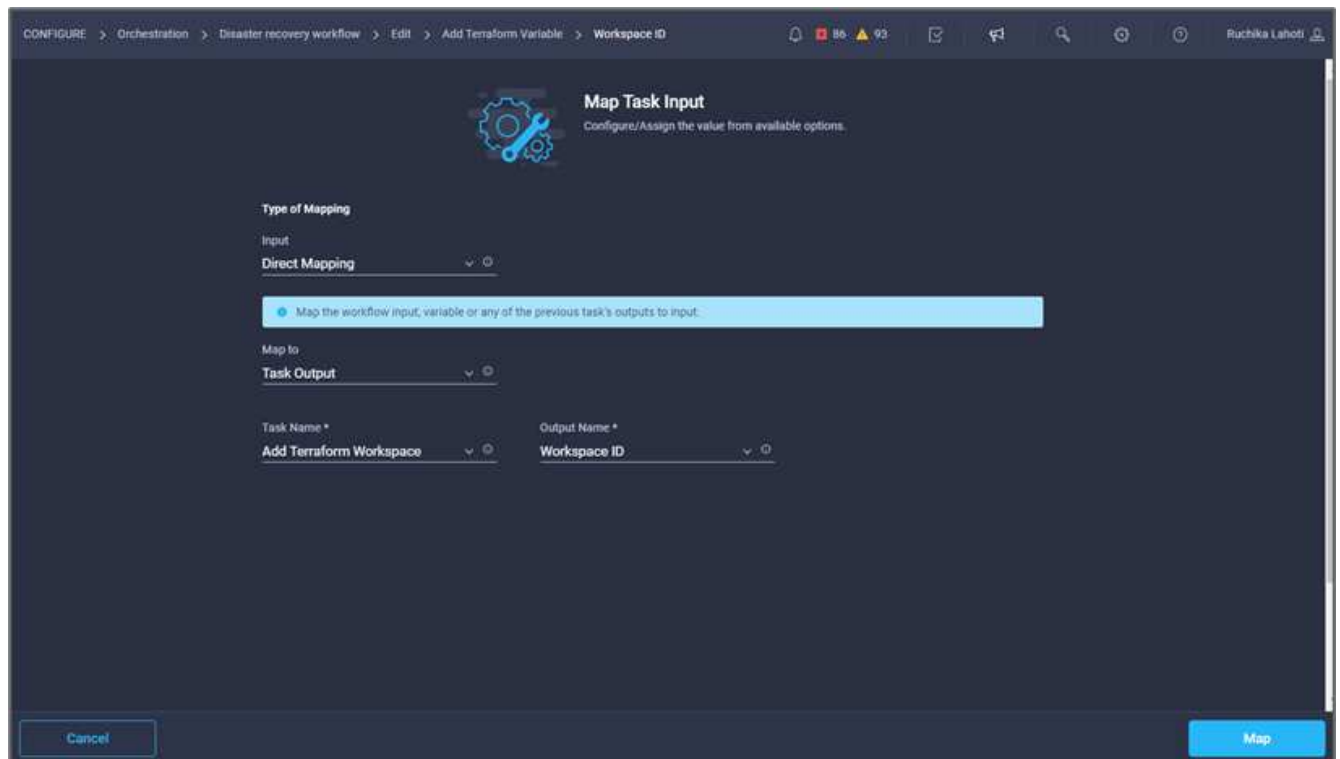
- Click **Add Terraform Variables**. In the **Workflow Properties** area, click the **General** tab. Optionally, you can change the name and description for this task.




- In the **Workflow Properties** area, click **Inputs**.
- Click **Map** in the **Terraform Cloud Target** field.
- Choose **Static Value** and click **Select Terraform Cloud Target**. Select the Terraform Cloud for Business account that was added as explained in [Configure Cisco Intersight Service for HashiCorp Terraform.](#)



- Click **Map**.
- Click **Map** in the ***Terraform Organization Name *** field.
- Choose **Static Value** and click **Select Terraform Organization**. Select the name of the Terraform Organization that you are part of in your Terraform Cloud for Business account.



11. Click **Map**.
12. Click **Map** in the **Terraform Workspace Name** field.
13. Choose **Direct Mapping** and click **Task Output**.
14. Click **Task Name** and click **Add Terraform Workspace**.



15. Click **Output Name** and click **Workspace Name**.

16. Click **Map**.
17. Click **Map** in the **Add Variables Options** field.
18. Choose **Direct Mapping** and click **Workflow Input**.
19. Click **Input Name** and **Create Workflow Input**.

Add Workflow Input

Display Name *
Terraform Variable

Reference Name *
TerraformAddVariable

Description
Terraform Variable to be added

Value Restrictions

☒ Required

☐ Collection/Multiple

Type
String

Min
0

Max
0

Regex

☐ Secure

☐ Object Selector

Cancel Add

20. In the Add Input wizard, complete the following steps:
 - a. Provide a display name and reference name (Optional).
 - b. Make sure to select **String** for the **Type**.
 - c. Click **Set Default Value and Override**.
 - d. Click **Variable Type** and then click **Non-Sensitive Variables**.

21. In the **Add Terraform Variables** section, provide the following information:

- **Key.** name_of_on-prem-ontap
- **Value.** Provide the name of on-premises ONTAP.
- **Description.** Name of the on-premises ONTAP.

22. Click + to add additional variables.

☒ Set Default Value ⓘ

☒ Allow User Override ⓘ

Default Values *

Terraform Variable

Key *

name_of_on-prem-ontap ⓘ

Value

Provide the name of On-premise ONTAP added in section Deploying ⓘ

Description

Name of the On-premise ONTAP ⓘ

☐ HCL ⓘ

+

Cancel Add

23. Add all the Terraform Variables as shown in the following table. You can also provide a default value.

Terraform variable name	Description
name_of_on-prem-ontap	Name of the on-premises ONTAP (FlexPod)

Terraform variable name	Description
on-prem-ontap_cluster_ip	The IP address of the storage cluster management interface
on-prem-ontap_user_name	Admin username for the storage cluster
Zone	GCP region where the working environment will be created
subnet_id	GCP subnet id where the working environment will be created
vpc_id	The VPC ID where the working environment will be created
capacity_package_name	The type of license to use
source_volume	The name of the source volume
source_storage_vm_name	The name of the source SVM
destination_volume	Name of volume on Cloud Volumes ONTAP
schedule_of_replication	The default is 1 hour
name_of_volume_to_create_on_cvo	Name of the cloud volume
workspace_id	The workspace_id where the working environment will be created
Project_id	The project_id where the working environment will be created
name_of_cvo_cluster	The name of the Cloud Volumes ONTAP working environment
gcp_service_account	gcp_service_account of Cloud Volumes ONTAP working environment

24. Click **Map** and then **Save**.

Add Terraform Variable

General

Inputs

Outputs

Variables

Search

Terraform Cloud Target *

Custom Value

Edit Mapping

View Value

Workspace ID *

Task Output

WorkspaceId | Add Terraform Work...

Edit Mapping

Terraform Variable

Workflow Input

Terraform Variables

Edit Mapping

Last saved an hour ago

Save

Execute

This completes the task of adding the required Terraform variables to the workspace. Next, add the required sensitive Terraform variables to the workspace. You can also combine both into a single task.

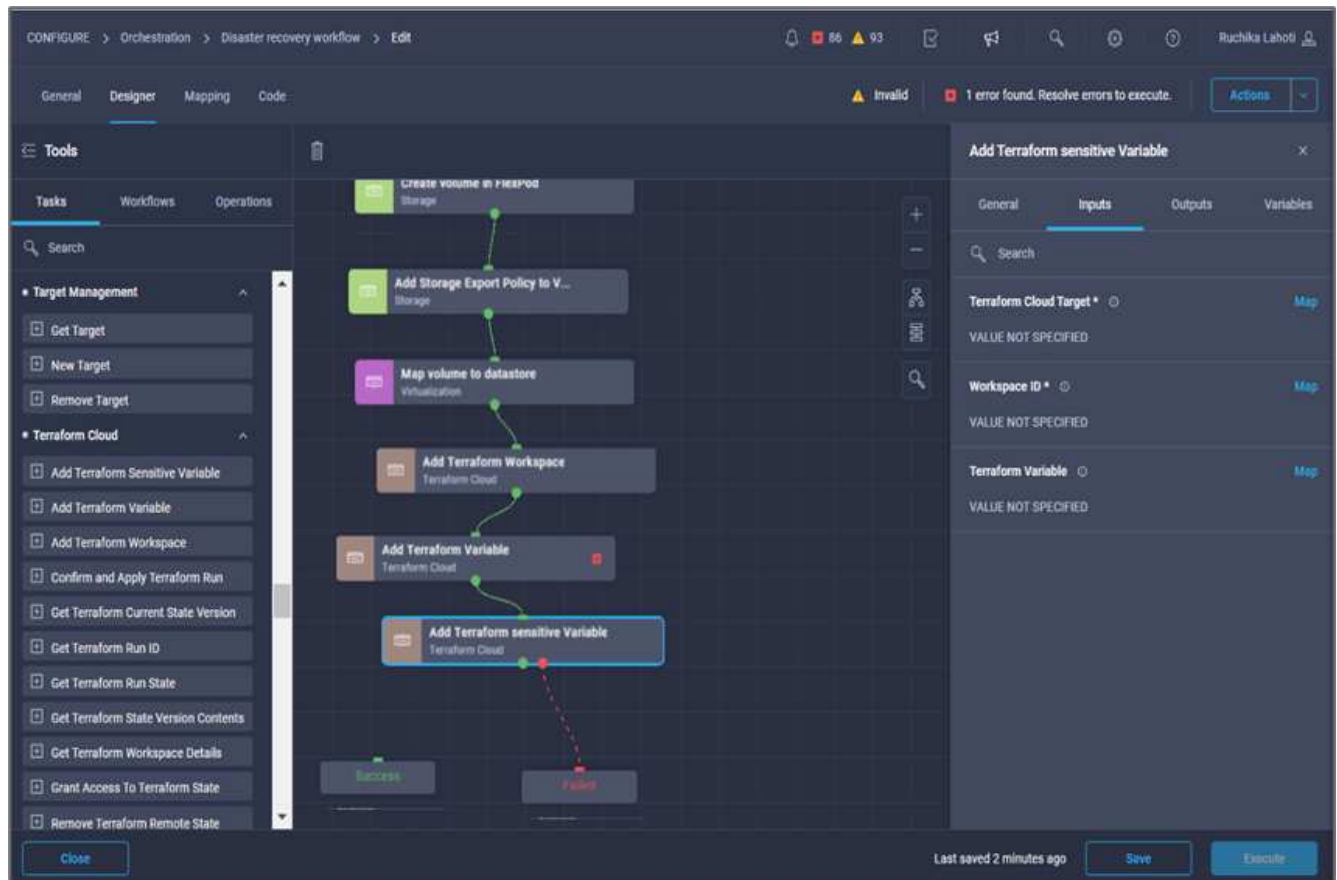
Procedure 7: Add sensitive variables to a workspace

1. Go to the **Designer** tab and click **Workflows** from the **Tools** section.
2. Drag and drop the **Terraform > Add Terraform Variables** workflow from the **Tools** section in the **Design** area.
3. Use Connector to connect the two **Add Terraform Workspace** tasks. Click **Save**.



A warning appears indicating that the two tasks have the same name. Ignore the error for now because you change the task name in the next step.

4. Click **Add Terraform Variables**. In the **Workflow Properties** area, click the **General** tab. Change the name to **Add Terraform Sensitive Variables**.



5. In the **Workflow Properties** area, click **Inputs**.
6. Click **Map** in the **Terraform Cloud Target** field.
7. Choose **Static Value** and click **Select Terraform Cloud Target**. Select the Terraform Cloud for Business account that was added in the section [Configure Cisco Intersight Service for HashiCorp Terraform.](#)
8. Click **Map**.
9. Click **Map** in the **Terraform Organization Name** field.
10. Choose **Static Value** and click **Select Terraform Organization**. Select the name of the Terraform Organization that you are part of in your Terraform Cloud for Business account.
11. Click **Map**.
12. Click **Map** in the **Terraform Workspace Name** field.

13. Choose **Direct Mapping** and click **Task Output**.
14. Click **Task Name** and then click **Add Terraform Workspace**.
15. Click **Output Name** and click the output **Workspace Name**.
16. Click **Map**.
17. Click **Map** in the **Add Variables Options** field.
18. Choose **Direct Mapping** and then click **Workflow Input**.
19. Click **Input Name** and **Create Workflow Input**.
20. In the Add Input wizard, complete the following steps:
 - a. Provide a display name and reference name (optional).
 - b. Make sure to select **Terraform Add Variables Options** for the type.
 - c. Click **Set Default Value**.
 - d. Click **Variable Type** and then click **Sensitive Variables**.
 - e. Click **Add**.

Add Workflow Input

Display Name *
terraform sensitive variable ⓘ

Reference Name *
terraformensitivevariable ⓘ

Description
Add Variables ⓘ

Value Restrictions

☒ Required ⓘ

☐ Collection/Multiple ⓘ

Type
Terraform Add Variables Option ▼ ⓘ

☒ Set Default Value ⓘ

☐ Allow User Override ⓘ

Default Values *
terraform sensitive variable

Variable Type *
Sensitive Variables × ▼ ⓘ

Cancel Add

21. In the **Add Terraform Variables** section, provide the following information:

- **Key.** cloudmanager_refresh_token.
- **Value.** Input the refresh token for NetApp Cloud Manager API operations.
- **Description.** Refresh token.



For more information about obtaining a refresh token for the NetApp Cloud Manager API operations, see the section [“Set up environment prerequisites.”](#)

×

Add Workflow Input

☒ Set Default Value ⓘ

☐ Allow User Override ⓘ

Default Values *

terraform sensitive variable

Variable Type *

Sensitive Variables × ∨ ⓘ

Add Sensitive Terraform Variables

Key *

cloudmanager_refresh_token ⓘ

Value

ⓘ ⓘ

Description

cloudmanager refresh token ⓘ

☐ HCL ⓘ

+

Cancel

Add

22. Add all the Terraform sensitive variables as shown in the table below. You can also provide a default value.

Terraform sensitive variable name	Description
cloudmanager_refresh_token	Refresh token. Obtain it from:
connector_id	The client ID of the Cloud Manager Connector. Obtain it from
cvo_admin_password	The admin password for Cloud Volumes ONTAP

Terraform sensitive variable name	Description
on-prem-ontap_user_password	Admin password for the storage cluster

- Click **Map**. This completes the task of adding the required Terraform sensitive variables to workspace. Next, start a new Terraform plan in the configured workspace.

Procedure 8: Start a new Terraform plan

- Go to the **Designer** tab and click **Tasks** from the **Tools** section.
- Drag and drop the **Terraform Cloud > Start New Terraform Plan** task from the **Tools** section on the **Design** area.
- Use Connector to connect between the tasks **Add Terraform Sensitive Variables** and **Start New Terraform Plan** tasks. Click **Save**.
- Click **Start New Terraform Plan**. In the **Task Properties** area, click the **General** tab. Optionally, you can change the name and description for this task.

The screenshot displays the Cisco Intersight Designer interface. On the left, the 'Tools' panel shows a list of tasks, with 'Start New Terraform Plan' highlighted. The main workspace shows a workflow diagram with several tasks connected by arrows. The 'Start New Terraform Plan' task is highlighted with a blue border. On the right, the 'Task Properties' panel for 'Start New Terraform Plan' is open, showing the 'General' tab. The 'Name' field is set to 'Start New Terraform Plan', the 'Version' is '1 (default)', and the 'Task Type' is 'Start New Terraform Plan'. The 'User Description' field contains the text: 'Starts a new plan or destroys a plan in the given Terraform Workspace'. At the bottom right, there are buttons for 'Save' and 'Execute'.

- In the **Task Properties** area, click **Inputs**.
- Click **Map** in the **Terraform Cloud Target** field.
- Choose **Static Value** and click **Select Terraform Cloud Target**. Select the Terraform Cloud for Business account that was added in the section “Configuring Cisco Intersight Service for HashiCorp Terraform.”
- Click **Map**.

9. Click **Map** in the **Workspace ID** field.
10. Choose **Direct Mapping** and click **Task Output**.
11. Click **Task Name** and then click **Add Terraform Workspace**.

Map Task Input
Configure/Assign the value from available options.

Type of Mapping
Input
Direct Mapping

Map the workflow input, variable or any of the previous task's outputs to input.

Map to
Task Output

Task Name *
Add Terraform Workspace

Output Name *

Map

12. Click **Output Name**, **Workspace ID**, and then **Map**.
13. Click **Map** in the **Reason for starting plan** field.
14. Choose **Direct Mapping** and then click **Workflow Input**.
15. Click **Input Name** and then **Create Workflow Input**.
16. In the Add Input wizard, complete the following steps:
 - a. Provide a display name and reference name (optional).
 - b. Make sure to select **String** for the **Type**.
 - c. Click **Set Default Value and Override**.
 - d. Input a default value for **Reason for starting plan** and click **Add**.

Add Workflow Input

☒ Required ⓘ

☐ Collection/Multiple ⓘ

Type
String ▼ ⓘ

Min **0** ⓘ Max **0** ⓘ Regex ⓘ

☐ Secure ⓘ

☐ Object Selector ⓘ

☒ Set Default Value ⓘ

☒ Allow User Override ⓘ


Default Values *

*Reason for starting plan **

terraform plan for replication between onprem volume and CVO ⓘ

Cancel Add

17. Click **Map**.
18. Click **Map** in the **Plan Operation** field.
19. Choose **Static Value** and click **Plan Operation**. Click **new plan**.



20. Click **Map**.

21. Click **Save**.

This completes the task of adding a Terraform Plan in Terraform Cloud for Business account. Next, create a sleep task for a few seconds.

Procedure 9: Sleep task for synchronization

Terraform Apply requires RunID, which is generated as a part of the Terraform Plan task. Waiting a few seconds between the Terraform Plan and Terraform Apply actions avoids timing issues.

1. Go to the **Designer** tab and click **Tasks** from the **Tools** section.
2. Drag and drop the **Core Tasks > Sleep Task** from the **Tools** section in the **Design** area.
3. Use Connector to connect the tasks **Start New Terraform Plan** and **Sleep Task**. Click **Save**.

The screenshot shows the Terraform Cloud workflow editor. A sequence of tasks is arranged in a flow: Start (grey), Create volume in FlexPod (Storage, green), Add Storage Export Policy to V... (Storage, green), Map volume to datastore (Virtualization, purple), Add Terraform Workspace (Terraform Cloud, orange), Add Terraform Variable (Terraform Cloud, orange), Add Terraform sensitive Variable (Terraform Cloud, orange), Start New Terraform Plan (Terraform Cloud, orange), Sleep Task (CoreTasks, blue), Success (grey), and Failed (grey). The Sleep Task is selected, and its properties are shown in a sidebar on the right.

Sleep Task Properties:

- Name ***: Sleep Task
- Version**: 1 (default)
- Task Type**: Sleep Task
- User Description**: Pauses the current workflow for the specified duration.
- Task Details**: Pauses the current workflow for the specified duration.

Buttons at the bottom: Last saved 2 minutes ago, Save, Execute.

- Click **Sleep Task**. In the **Task Properties** area, click the **General** tab. Optionally, you can change the name and description for this task. In this example, the name of the task is **Synchronize**.
- In the **Task Properties** area, click **Inputs**.
- Click **Map** in the **Sleep Time in Seconds** field.
- Choose **Static Value** and input **15** in for the **Sleep Time in Seconds**.

The screenshot shows the 'Edit Task Input Mapping' dialog. It has a title 'Edit Task Input Mapping' and a subtitle 'Configure/Assign the value from available options.' Below the title is a gear icon. The 'Type of Mapping' is set to 'Static Value'. The 'Input' field is set to 'Static Value'. The 'Sleep Time in Seconds' field is set to '15'. There is a range indicator '1 - 600' below the field.

Edit Task Input Mapping
Configure/Assign the value from available options.

Type of Mapping
Input
Static Value

Provide custom values as the input.

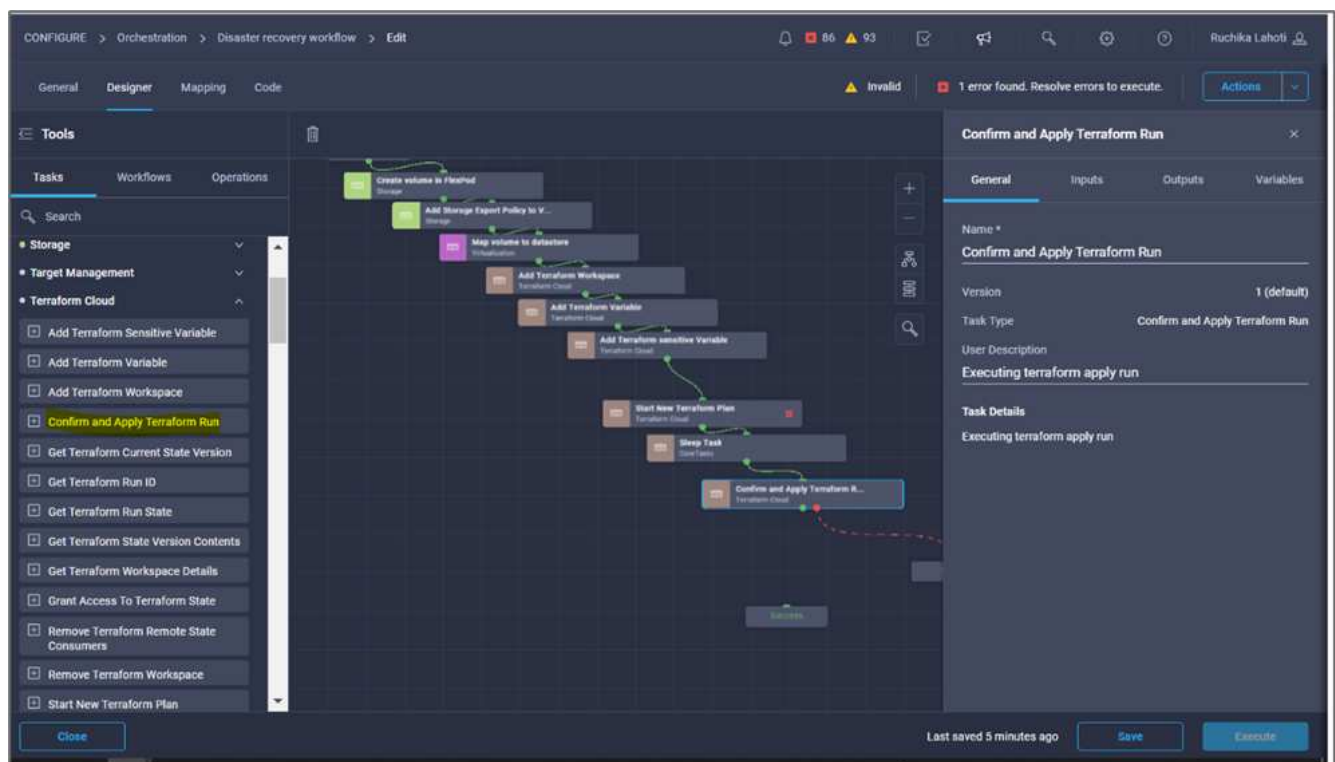
Sleep Time in Seconds *
15
1 - 600

8. Click **Map**.
9. Click **Save**.

This completes the sleep task. Next, create the last task of this workflow, confirming and applying the Terraform Run.

Procedure 10: Confirm and apply Terraform Run

1. Go to the **Designer** tab and click **Tasks** from the **Tools** section.
2. Drag and drop the **Terraform Cloud > Confirm and Apply Terraform Run** task from the **Tools** section in the **Design** area.
3. Use connector to connect the tasks **Synchronize** and **Confirm and Apply Terraform Run**. Click **Save**.
4. Click **Confirm** and **Apply Terraform Run**. In the **Task Properties** area, click the **General** tab. Optionally, you can change the name and description for this task.



5. In the **Task Properties** area, click **Inputs**.
6. Click **Map** in the **Terraform Cloud Target** field.
7. Choose **Static Value** and click **Select Terraform Cloud Target**. Select the Terraform Cloud for Business account that was added in [Configure Cisco Intersight Service for HashiCorp Terraform.](#)
8. Click **Map**.
9. Click **Map** in the **Run ID** field.
10. Choose **Direct Mapping** and click **Task Output**.
11. Click **Task Name** and click **Start New Terraform Plan**.
12. Click **Output Name** and then click **Run ID**.

CONFIGURE > Orchestration > Disaster recovery workflow > Edit > Confirm and Apply Terraform Run > Run ID

86 93

Ruchika Lahoti

Map Task Input

Configure/Assign the value from available options.

Type of Mapping

Input
Direct Mapping

Map the workflow input, variable or any of the previous task's outputs to input.

Map to

Task Output

Task Name *
Start New Terraform Plan

Output Name *
Run ID

Cancel Map

13. Click **Map**.
14. Click **Save**.
15. Click **Auto Align Workflow** so that all tasks are aligned. Click **Save**.



This completes the Confirm and Apply Terraform Run task. Use Connector to connect between the **Confirm and Apply Terraform Run** task and the **Success** and **Failed** tasks.

Procedure 11: Import a Cisco-built workflow

Cisco Intersight Cloud Orchestrator enables you to export workflows from a Cisco Intersight account to your system and then import them to another account. A JSON file was created by exporting the built workflow that can be imported to your account.

A JSON file for the workflow component is available in the [GitHub repository](#).

Next: Terraform execution from controller.

Terraform execution from controller

Previous: DR workflow.

We can execute the Terraform plan using a controller. You can skip this section if you have already executed your Terraform plan using an ICO workflow.

Prerequisites

Setup of the solution begins with a management workstation that has access to the Internet and with a working installation of Terraform.

A guide for installing Terraform can be found [here](#).

Clone GitHub repo

The first step in the process is to clone the GitHub repo to a new empty folder on the management workstation. To clone the GitHub repository, complete the following steps:

1. From the management workstation, create a new folder for the project. Create a new folder inside this folder named `/root/snapmirror-cvo` and Clone the GitHub repo into it.
2. Open a command-line or console interface on the management workstation and change directories to the new folder just created.
3. Clone the GitHub collection using the following command:

```
Git clone https://github.com/NetApp-Automation/FlexPod-hybrid-cloud-for-GCP-with-Intersight-and-CVO
```

1. Change directories to the new folder named `snapmirror-cvo`.

Terraform execution



- **Init.** Initialize the (local) Terraform environment. Usually executed only once per session.
- **Plan.** Compare the terraform state with the as-in state in the cloud and build and display an execution plan. This does not change the deployment (read-only).
- **Apply.** Apply the plan from the plan phase. This potentially changes the deployment (read and write).
- **Destroy.** All resources that are governed by this specific terraform environment.

For details, see [here](#).

Solution validation

In this section, we revisit the solution with a sample data-replication workflow and take a few measurements to verify the integrity of data replication from the NetApp ONTAP instance running in FlexPod to NetApp Cloud Volumes ONTAP running on Google Cloud.

We used the Cisco Intersight workflow orchestrator in this solution and will continue to use this for our use case.

Notably, the limited set of Cisco Intersight workflows used in this solution do not represent the full set of workflows that Cisco Intersight is equipped with. You can create custom workflows based on your specific requirements and have them triggered from Cisco Intersight.

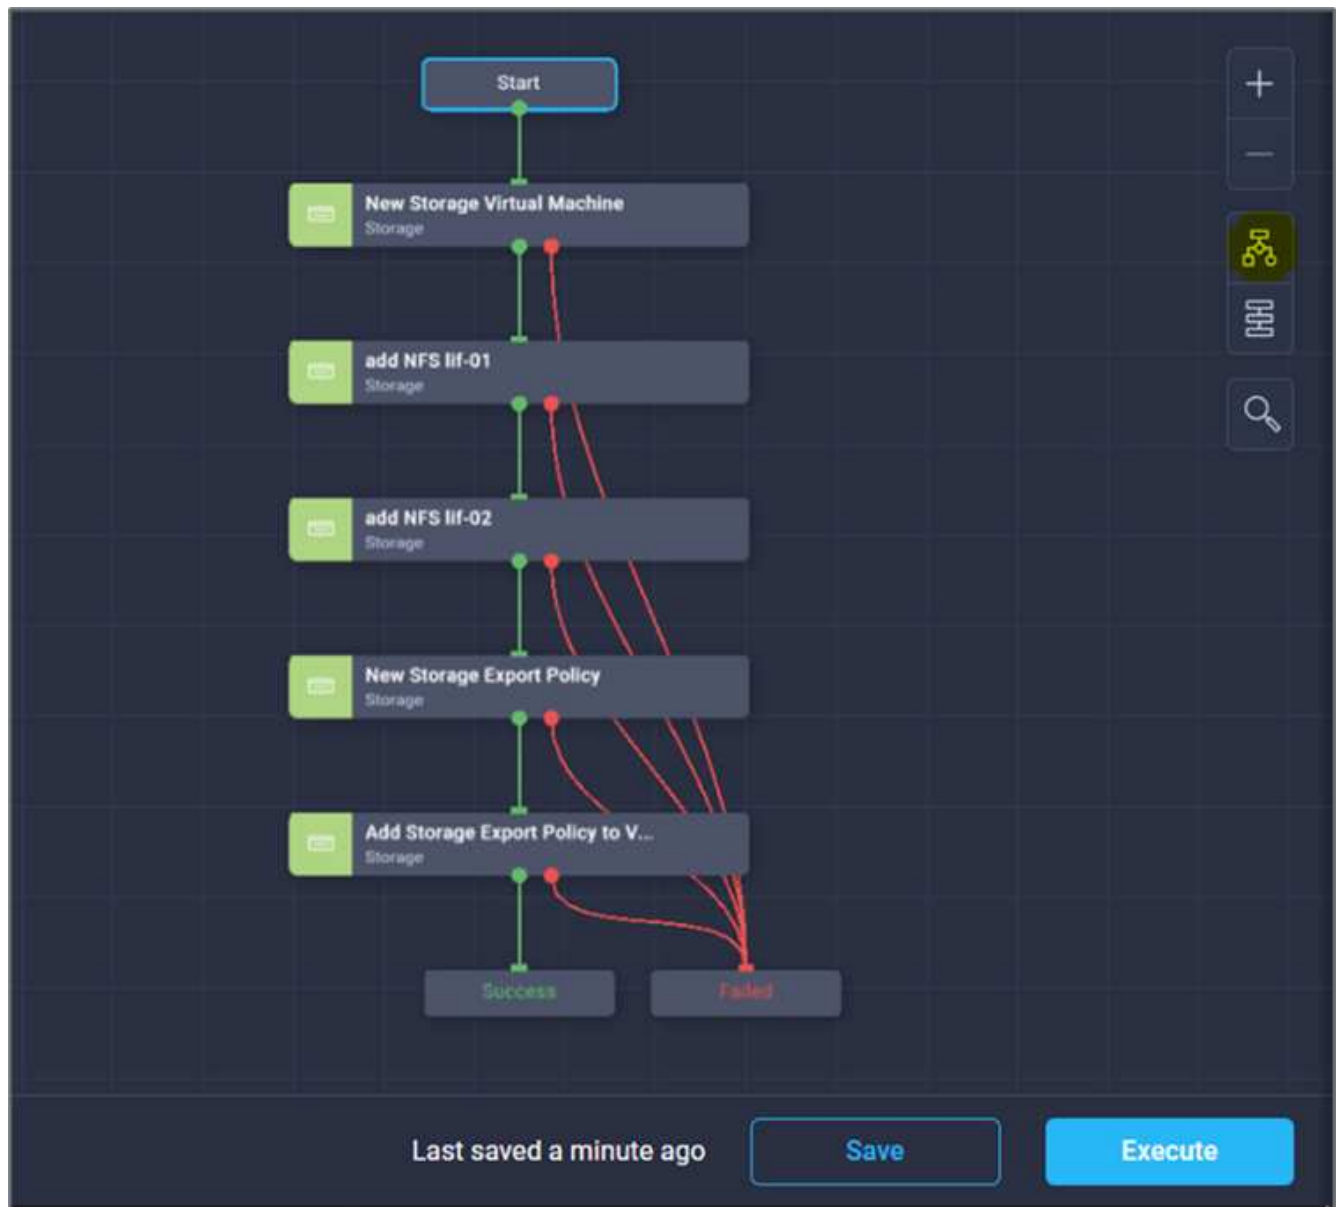
To perform the validation of a successful DR scenario, first move data from a volume in ONTAP that is part of FlexPod to Cloud Volumes ONTAP using SnapMirror. Then you can attempt to access the data from the Google cloud compute instance followed by a data integrity check.

The following high-level steps are used to verify the success criteria of this solution:

1. Generate an SHA256 checksum on the sample dataset that is present in an ONTAP volume in FlexPod.
2. Set up a volume SnapMirror relationship between ONTAP in FlexPod and Cloud Volumes ONTAP.
3. Replicate the sample dataset from FlexPod to Cloud Volumes ONTAP.
4. Break the SnapMirror relationship and promote the volume in Cloud Volumes ONTAP to production.
5. Map the Cloud Volumes ONTAP volume with the dataset to a compute instance in Google Cloud.
6. Generate an SHA256 checksum on the sample dataset in Cloud Volumes ONTAP.
7. Compare the checksum on the source and destination; presumably, the checksums on both sides match.

To execute the on-premises workflow, complete the following steps:

1. Create a workflow in Intersight for on-premises FlexPod.



2. Provide the required inputs and execute the workflow.

Execute Workflow: Configure on-prem FlexPod storage

Execute Workflow
Fill Attributes

General

Organization *
default

Workflow Instance Name
Configure on-prem FlexPod storage

Workflow Inputs

Storage Virtual Machine *
flexpod-svm

Storage Vendor Virtual Machine Options

Platform Type
☐ Pure FlashArray
 ☐ Hitachi Virtual Storage Platform
 ☒ NetApp Active IQ Unified Manager
 ☐ None

NetApp Virtual Machine Options

Storage VM Protocols *
NFS

Storage VM Protocols *
iSCSI

☐ Manage Administrator Account: vsadmin

Route Destination IPv4 Gateway
10.61.183.1

Execute

3. Verify the newly created SVM in the system manager.

ONTAP System Manager Search actions, objects, and pages

DASHBOARD

INSIGHTS

STORAGE

Overview

Volumes

LUNs

Consistency Groups

NVMe Namespaces

Shares

Qtrees

Quotas

Storage VMs

Tiers

Storage VMs

+ Add More

Name
flexpod-svm
hybrid-cloud-svm
hybrid_cloud_2_svm
infra_svm
nvme1
terraform-demo-svm

flexpod-svm All Storage VMs

Overview Settings Snap

Security

Certificates

4. Create and execute another disaster recovery workflow to create a volume in on-prem FlexPod and establish a SnapMirror relationship between this volume in FlexPod and Cloud Volumes ONTAP.



5. Verify the newly created volume in ONTAP system manager.

ONTAP System Manager

Search actions, objects, and pages

DASHBOARD

INSIGHTS

STORAGE

Overview

Volumes

LUNs

Consistency Groups

NVMe Namespaces

Shares

Qtrees

Quotas

Storage VMs

Tiers

Volumes

+ Add

More

	Name	Storage VM	Status	Capacity
		hybrid-cloud-svr	(All)	>
	application_copy	hybrid-cloud-svm	Online	3.12 MiB used 19 GiB available 20 GiB
	audit_log_vol	hybrid-cloud-svm	Online	32.7 MiB used 200 GiB available 200 GiB
	hybrid_cloud_svm_root	hybrid-cloud-svm	Online	1.68 MiB used 971 MiB available 1 GiB
	test	hybrid-cloud-svm	Online	648 KiB used 972 MiB available 1 GiB
	Test_Vol1	hybrid-cloud-svm	Online	10.6 MiB used 9.99 GiB available 10 GiB

- Mount the same NFS volume to an on-premises virtual machine, then copy the sample dataset and perform the checksum.

```

root@hybridcloudbackup:/snapmirror_demo# mount -t nfs 172.22.4.157:/Test_Vol1 /snapmirror_demo
root@hybridcloudbackup:/snapmirror_demo# df -kh
Filesystem      Size  Used Avail Use% Mounted on
udev            1.9G   0    1.9G   0% /dev
tmpfs           394M  1.1M  393M   1% /run
/dev/sda2       16G   11G   4.2G  72% /
tmpfs           2.0G   0    2.0G   0% /dev/shm
tmpfs           5.0M   0    5.0M   0% /run/lock
tmpfs           2.0G   0    2.0G   0% /sys/fs/cgroup
/dev/loop1      55M   55M    0 100% /snap/core18/1705
/dev/loop2      69M   69M    0 100% /snap/lxd/14804
/dev/loop0      28M   28M    0 100% /snap/snapd/7264
172.22.4.157:/Test_Vol1 10G 512K 10G   1% /snapmirror_demo
root@hybridcloudbackup:/snapmirror_demo#


```

```

root@hybridcloudbackup:/snapmirror_demo#
root@hybridcloudbackup:/snapmirror_demo# sha256sum test.zip
888a23c8495ad33fdf11a931ffc344c3643f15d5cefedbbf1326016e31ec5a59 test.zip
root@hybridcloudbackup:/snapmirror_demo#
root@hybridcloudbackup:/snapmirror_demo#


```

- Check the replication status in Cloud Manager. The data transfer can take few minutes based on the size of the data. After it is completed, you can see the SnapMirror status as **Idle**.


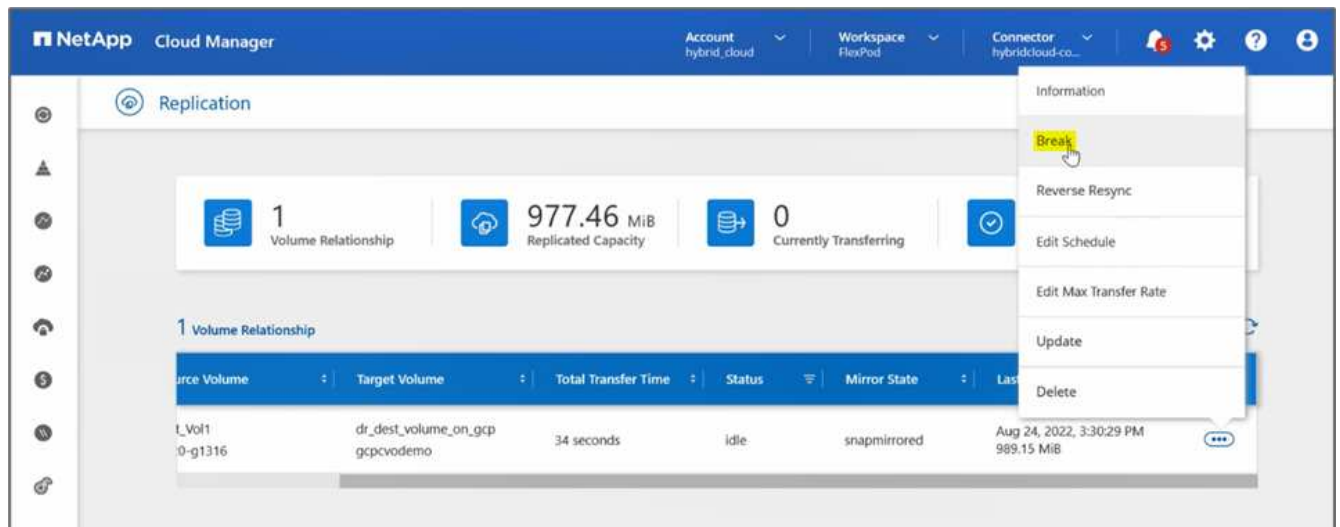


8. When the data transfer is complete, simulate a disaster on the source side by stopping the SVM that hosts the `Test_vol1` volume.

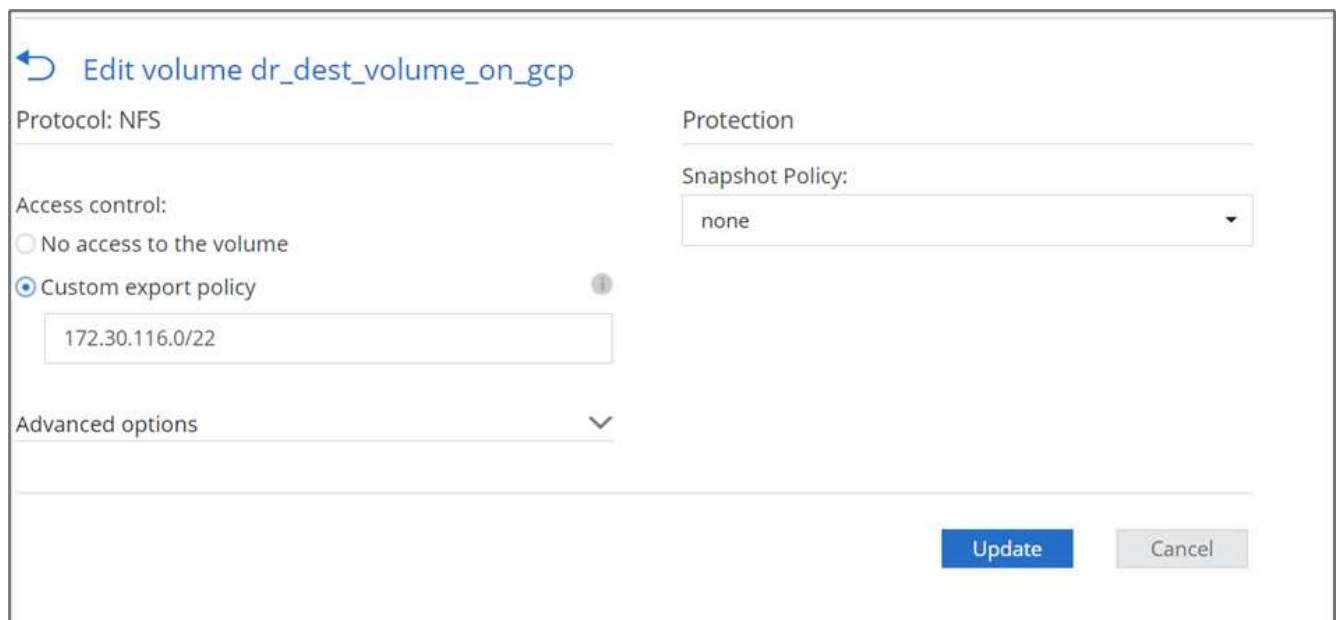
After the SVM has been stopped, the `Test_vol1` volume is not visible in the Cloud Manager.




9. Break the replication relationship and promote the Cloud Volumes ONTAP destination volume to production.



10. Edit the volume and enable client access by associating it with an export policy.



11. Obtain the ready-to-use mount command for the volume.



Mount Volume dr_dest_volume_on_gcp

Go to your Linux machine and enter this mount command

```
mount 172.30.116.153:/dr_dest_volume_on_gcp <dest...
```

[Copy](#)

12. Mount the volume to a compute instance, verify that the data is present in the destination volume, and generate the SHA256 checksum of the `sample_dataset_2GB` file.

```
drwxr-xr-x 21 root    root          4096 Aug 24 10:20 ../
-rwxr-xr-x  1 nobody 4294967294 1015306240 Aug 24 09:59 test.zip*
ruchikal_netapp_com@demo-nfs:/snapmirror_dest$
ruchikal_netapp_com@demo-nfs:/snapmirror_dest$ sha256sum test.zip
888a23c8495ad33fdf11a931ffc344c3643f15d5cefedbbf1326016e31ec5a59 test.zip
ruchikal_netapp_com@demo-nfs:/snapmirror_dest$
```

13. Compare the checksum values at both the source (FlexPod) and the destination (Cloud Volumes ONTAP).
14. The checksums match to the source and destination.

You can confirm that the data replication from the source to the destination was completed successfully and the data integrity was maintained. This data can now be safely consumed by the applications to serve clients while the source site goes through restoration.

Next: [Conclusion](#).

Conclusion

[Previous: Solution validation.](#)

In this solution, the NetApp Cloud Data service, Cloud Volumes ONTAP, and FlexPod Datacenter infrastructure were used to build a DR solution with a public cloud powered by the Cisco Intersight Cloud Orchestrator. The FlexPod solution has constantly evolved to enable customers to modernize their applications and business-delivery processes. With this solution, you can build a BCDR plan with the public cloud as your go-to location for a transient or full-time DR plan while keeping the cost of the DR solution low.

Data replication between on-premises FlexPod and NetApp Cloud Volumes ONTAP was handled by proven SnapMirror technology, but you can also select other NetApp data-transfer and synchronization tools like Cloud Sync for your data mobility requirements. Security of the data in-flight provided by built-in encryption technologies based on TLS/AES.

Whether you have a temporary DR plan for an application or a full-time DR plan for a business, the portfolio of products used in this solution can meet both requirements at scale. Powered by Cisco Intersight Workflow Orchestrator, the same can be automated with prebuilt workflows that not just eliminate the need to rebuild processes but also accelerate the implementation of a BCDR plan.

The solution enables the management of FlexPod on-premises and data replication across a hybrid cloud in a very easy and convenient manner with automation and orchestration provided by Cisco Intersight Cloud Orchestrator.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

GitHub

- All Terraform Configurations used

<https://github.com/NetApp-Automation/FlexPod-hybrid-cloud-for-GCP-with-Intersight-and-CVO>

- JSON files for importing workflows

https://github.com/ucs-compute-solutions/FlexPod_DR_Workflows

Cisco Intersight

- Cisco Intersight Help Center

<https://intersight.com/help/saas/home>

- Cisco Intersight Cloud Orchestrator Documentation:

https://intersight.com/help/saas/features/orchestration/configure#intersight_cloud_orchestrator

- Cisco Intersight Service for HashiCorp Terraform Documentation

https://intersight.com/help/saas/features/terraform_cloud/admin

- Cisco Intersight Data Sheet

<https://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/intersight/intersight-ds.html>

- Cisco Intersight Cloud Orchestrator Data Sheet

<https://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/intersight/nb-06-intersight-cloud-orch-aag-cte-en.html>

- Cisco Intersight Service for HashiCorp Terraform Data Sheet

<https://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/intersight/nb-06-intersight-terraf-ser-aag-cte-en.html>

FlexPod

- FlexPod Home Page

<https://www.flexpod.com>

- Cisco Validated Design and deployment guides for FlexPod

[FlexPod Datacenter with Cisco UCS 4.2\(1\) in UCS Managed Mode, VMware vSphere 7.0 U2, and NetApp ONTAP 9.9 Design Guide](#)

- FlexPod Datacenter with Cisco UCS X-Series

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_xseries_esxi7u2_design.html

Interoperability

- NetApp Interoperability Matrix Tool

<http://support.netapp.com/matrix/>

- Cisco UCS Hardware and Software Interoperability Tool

<http://www.cisco.com/web/techdoc/ucs/interoperability/matrix/matrix.html>

- VMware Compatibility Guide

<http://www.vmware.com/resources/compatibility/search.php>

NetApp Cloud Volumes ONTAP reference documents

- NetApp Cloud Manager

https://docs.netapp.com/us-en/occm/concept_overview.html

- Cloud Volumes ONTAP

<https://docs.netapp.com/us-en/cloud-manager-cloud-volumes-ontap/task-getting-started-gcp.html>

- Cloud Volumes ONTAP TCO Calculator

<https://cloud.netapp.com/google-cloud-calculator>

- Cloud Volumes ONTAP Sizer

<https://cloud.netapp.com/cvo-sizer>

- Cloud Assessment Tool

<https://cloud.netapp.com/assessments>

- NetApp Hybrid Cloud

<https://cloud.netapp.com/hybrid-cloud>

- Cloud Manager API documentation

https://docs.netapp.com/us-en/occm/reference_infrastructure_as_code.html

Troubleshooting issues

[https://kb.netapp.com/Advice_and_Troubleshooting/Cloud_Services/Cloud_Volumes_ONTAP_\(CVO\)](https://kb.netapp.com/Advice_and_Troubleshooting/Cloud_Services/Cloud_Volumes_ONTAP_(CVO))

Terraform

- Terraform Cloud

<https://www.terraform.io/cloud>

- Terraform Documentation

<https://www.terraform.io/docs/>

- NetApp Cloud Manager Registry

<https://registry.terraform.io/providers/NetApp/netapp-cloudmanager/latest>

GCP

- ONTAP High Availability for GCP

<https://cloud.netapp.com/blog/gcp-cvo-blg-what-makes-cloud-volumes-ontap-high-availability-for-gcp-tick>

- GCP perquisite

<https://netapp.hosted.panopto.com/Panopto/Pages/Viewer.aspx?id=f3d0368b-7165-4d43-a76e-ae01011853d6>

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.